

Camera Algorithm and Tuning Guide

The Raspberry Pi Camera Algorithm and Tuning Guide

Colophon

© 2022-2026 Raspberry Pi Ltd

This documentation is licensed under a Creative Commons Attribution-NoDerivatives 4.0 International (CC BY-ND).

Release	2
Build date	20/10/2026
Build version	7301bc91433c

Legal disclaimer notice

TECHNICAL AND RELIABILITY DATA FOR RASPBERRY PI PRODUCTS (INCLUDING DATASHEETS) AS MODIFIED FROM TIME TO TIME ("RESOURCES") ARE PROVIDED BY RASPBERRY PI LTD ("RPL") "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. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW IN NO EVENT SHALL RPL 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 THE RESOURCES, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

RPL reserves the right to make any enhancements, improvements, corrections or any other modifications to the RESOURCES or any products described in them at any time and without further notice.

The RESOURCES are intended for skilled users with suitable levels of design knowledge. Users are solely responsible for their selection and use of the RESOURCES and any application of the products described in them. User agrees to indemnify and hold RPL harmless against all liabilities, costs, damages or other losses arising out of their use of the RESOURCES.

RPL grants users permission to use the RESOURCES solely in conjunction with the Raspberry Pi products. All other use of the RESOURCES is prohibited. No licence is granted to any other RPL or other third party intellectual property right.

HIGH RISK ACTIVITIES. Raspberry Pi products are not designed, manufactured or intended for use in hazardous environments requiring fail safe performance, such as in the operation of nuclear facilities, aircraft navigation or communication systems, air traffic control, weapons systems or safety-critical applications (including life support systems and other medical devices), in which the failure of the products could lead directly to death, personal injury or severe physical or environmental damage ("High Risk Activities"). RPL specifically disclaims any express or implied warranty of fitness for High Risk Activities and accepts no liability for use or inclusions of Raspberry Pi products in High Risk Activities.

Raspberry Pi products are provided subject to RPL's Standard Terms. RPL's provision of the RESOURCES does not expand or otherwise modify RPL's Standard Terms including but not limited to the disclaimers and warranties expressed in them.

Table of Contents

1. Introduction 5
2. Overview 7
    2.1. The Camera Module 7
    2.2. CSI Connector 7
    2.3. On-Chip Hardware 7
    2.4. Software and Control Algorithms 8
3. Driver Framework 9
    3.1. Camera Drivers 9
    3.2. Device Tree 9
    3.3. The `CamHelper` Class 9
4. Control Algorithm Overview 11
    4.1. Framework 11
    4.2. Defining and Loading Algorithms 11
    4.3. Standard Algorithms 11
    4.4. Algorithm Communication 12
    4.5. Performance Considerations 13
    4.6. Example Camera Tuning File 13
        4.6.1. Pi 4 and Earlier Devices 13
        4.6.2. Pi 5 Devices 15
    4.7. Directory Structure 17
5. Raspberry Pi Control Algorithms 19
    5.1. Black Level 19
        5.1.1. Name 19
        5.1.2. Overview 19
        5.1.3. Metadata Dependencies 19
    5.2. Defective Pixel Correction (DPC) 19
        5.2.1. Name 19
        5.2.2. Overview 20
        5.2.3. Parameters 20
        5.2.4. Metadata Dependencies 20
    5.3. Lux 20
        5.3.1. Name 20
        5.3.2. Overview 21
        5.3.3. Parameters 21
        5.3.4. Metadata Dependencies 21
    5.4. Noise 21
        5.4.1. Name 21
        5.4.2. Overview 21
        5.4.3. Parameters 21
        5.4.4. Metadata Dependencies 21
    5.5. GEQ (Green Equalisation) 22
        5.5.1. Name 22
        5.5.2. Overview 22
        5.5.3. Parameters 22
        5.5.4. Metadata Dependencies 23
    5.6. SDN (Spatial Denoise) 23
        5.6.1. Name 23
        5.6.2. Overview 23
        5.6.3. External API 24
        5.6.4. Parameters 24
        5.6.5. Metadata Dependencies 24
    5.7. Denoise (Combined Denoise Algorithms) 24
        5.7.1. Name 24
        5.7.2. Overview 24

5.7.3. External API 25
5.7.4. Parameters 25
5.7.5. Metadata Dependencies 26
5.8. AWB (Automatic White Balance) 26
5.8.1. Name 26
5.8.2. Overview 26
5.8.3. The Feasible Set 26
5.8.4. Bayes Theorem 27
5.8.5. The Priors 28
5.8.6. Calculating the Posteriors 28
5.8.7. Log Likelihood 29
5.8.8. Implementation 29
5.8.9. External API 29
5.8.10. Parameters 30
5.8.11. Metadata Dependencies 30
5.8.12. Extensions 31
5.8.13. Uniform Colour 31
5.8.14. Non-grey Colours 31
5.9. AGC/AEC (Automatic Gain Control / Automatic Exposure Control) 31
5.9.1. Name 31
5.9.2. Overview 31
5.9.3. Region-based Metering 32
5.9.4. Histogram Constraints 32
5.9.5. Channel Constraints 34
5.9.6. Exposure Modes 34
5.9.7. Implementation 35
5.9.8. External API 35
5.9.9. Parameters 36
5.9.10. Metadata Dependencies 37
5.10. ALSC (Automatic Lens Shading Correction) 37
5.10.1. Name 37
5.10.2. Overview 37
5.10.3. Luminance Shading 38
5.10.4. Colour Shading 39
5.10.5. Colour Temperature based calibration 39
5.10.6. Adaptive ALSC Algorithm 39
5.10.7. Solution 40
5.10.8. Implementation 40
5.10.9. Parameters 41
5.10.10. Metadata Dependencies 41
5.11. Contrast 41
5.11.1. Name 41
5.11.2. Overview 41
5.11.3. Adaptive Contrast Enhancement 42
5.11.4. External API 42
5.11.5. Parameters 43
5.11.6. Metadata Dependencies 43
5.12. CCM (Colour Correction Matrices) 43
5.12.1. Name 43
5.12.2. Overview 43
5.12.3. External API 44
5.12.4. Parameters 44
5.12.5. Metadata Dependencies 44
5.13. Sharpening 44
5.13.1. Name 44
5.13.2. Overview 45
5.13.3. External API 45
5.13.4. Parameters 45

5.14. Auto Focus 45
5.14.1. Name 45
5.14.2. Overview 45
5.14.3. External API 46
5.14.4. Parameters 46
5.14.5. Metadata Dependencies 47
5.15. Chromatic Aberration Correction 47
5.15.1. Name 47
5.15.2. Overview 47
5.15.3. External API 47
5.15.4. Parameters 48
5.16. HDR (High Dynamic Range) 48
5.16.1. Name 48
5.16.2. Overview 48
5.16.3. External API 49
5.16.4. Parameters 49
5.16.5. Metadata Dependencies 50
5.17. Software Camera Synchronisation 50
5.17.1. Name 50
5.17.2. Overview 50
5.17.3. Metadata Dependencies 51
5.18. Metadata and Statistics Usage 51
6. Camera Tuning Tool 52
6.1. Overview 52
6.2. Raspberry Pi *rpicam-apps* 52
6.3. Software Installation 52
6.4. Equipment 52
6.4.1. X-rite (Macbeth) Colour Checker 52
6.4.2. Colorimeter 53
6.4.3. Integrating Spheres and Flat Field LEDs 54
6.5. Capturing Calibration Images with *libcamera* 54
6.5.1. Capturing a Raw Image 54
6.6. Image Capture Requirements 55
6.6.1. Macbeth Chart Images 55
6.6.2. Lens Shading Images 56
6.6.3. Chromatic Aberration Correction (CAC) Images 57
6.7. Creating the Tuning 58
6.7.1. Collecting the Files 58
6.7.2. Running the Tool 58
6.8. Tuning only the Lens Shading Correction 59
6.9. Tweaking the Tuning produced by the Tool 59
6.9.1. Blocks not Tuned 59
6.9.2. Guidance on how to Tweak the Tuning 60
6.9.3. Black Level (`rpi.black_level`) 60
6.9.4. Defective Pixel Correction (`rpi.dpc`) 60
6.9.5. Lux (`rpi.lux`) 60
6.9.6. Noise (`rpi.noise`) 60
6.9.7. Green Equalisation (`rpi.geq`) 60
6.9.8. Spatial Denoise (`rpi.sdn` or `rpi.denoise`) 60
6.9.9. AWB (`rpi.awb`) 61
6.9.10. AGC/AEC (`rpi.agc`) 62
6.9.11. Automatic Lens Shading Correction (`rpi.alsc`) 63
6.9.12. Contrast (`rpi.contrast`) 63
6.9.13. Colour Correction Matrices (`rpi.ccm`) 64
6.9.14. Sharpening (`rpi.sharpen`) 64

1. Introduction

The Camera Algorithm and Tuning Guide is intended for users of the Raspberry Pi computer with an image sensor (camera) connected through the Raspberry Pi's CSI (Camera Serial Interface) camera port, such as one of the following.

- The v1 camera based on the Omnivision OV5647.
- The v2 camera based on the Sony imx219.
- The HQ Camera, based on the Sony imx477.
- The Camera Module 3, based on the Sony imx708.
- The Global Shutter Camera, based on the Sony imx296.

There are also a number of compatible camera modules available from third parties, though information and support for such modules will come from the respective third party vendor.

Figure 1.
Connecting a camera board through the CSI-2 port. Raspberry Pi 5 has a pair of mini CSI-2 connectors in the same location.



The software stack driving the camera system is *libcamera*. Experience has shown that driving complex camera systems directly through kernel (typically, V4L2) drivers is very difficult, often leading to large amounts of undesirable and highly platform-specific application code. For this reason a much higher level userspace camera stack, *libcamera*, has emerged providing mechanisms to integrate image sensors and Image Signal Processors (ISPs).

Here we describe just such an integration, showing how drivers can be written to make the components of the Raspberry Pi imaging system work with *libcamera*, concentrating in particular on the processes of calibrating and tuning the ISP to work well with different image sensors. Moreover, the *libcamera* integration avoids using any of the proprietary control algorithms supplied by chip vendors. Rather, Raspberry Pi is providing its own control algorithms, running directly on the chip's ARM cores, as open source code which can be easily inspected and modified by users.

This document is not specifically a tutorial on or guide to *libcamera* itself, for which the reader is referred to http://libcamera.org for more information.

The remainder of this document is organised as follows.

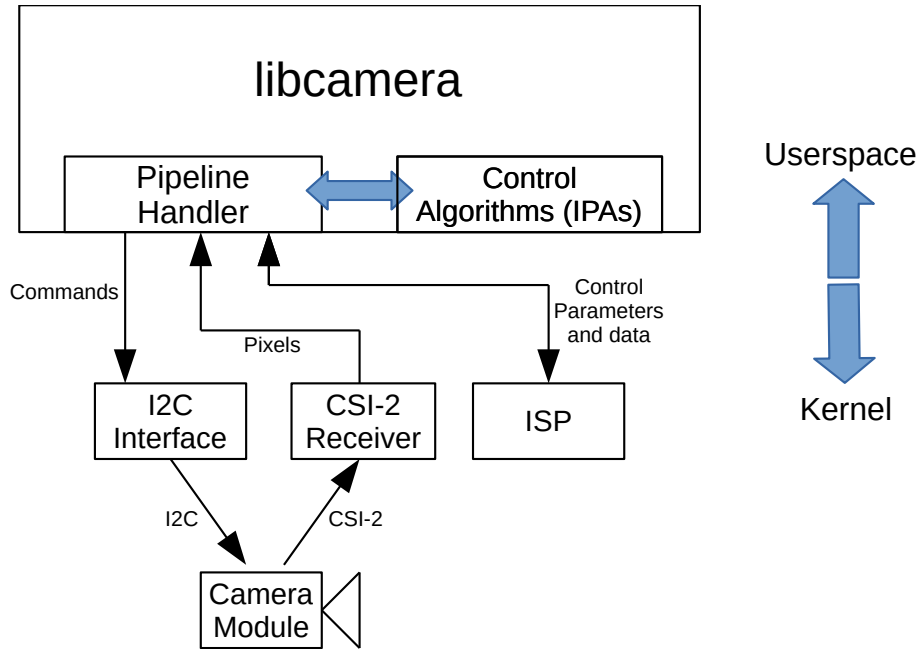
- Chapter 2: Overview. This is a brief overview of the system as a whole, showing what is already provided, and what must be added in terms of image sensor drivers and ISP control algorithms to obtain a successful images.
- Chapter 3: Camera Drivers. As far as possible the camera, or image sensor, is driven by a standard type of kernel driver. However, there are one or two details involved in making drivers work optimally in the Raspberry Pi framework.
- Chapter 4: Control Algorithm Overview. Raspberry Pi provides a control algorithm framework which makes it very easy for applications to use Raspberry Pi's, or anyone else's, ISP control algorithms.
- Chapter 5: Raspberry Pi Control Algorithms. Here we supply a detailed description of Raspberry Pi's implementation of the control algorithms.
- Chapter 6: Camera Tuning Tool. Finally, there is also a Camera Tuning Tool which automates most of the process of tuning the ISP control algorithms for different image sensors.

For further help beyond the contents of this user guide, please visit the Raspberry Pi Camera Board forum at https://forums.raspberrypi.com/viewforum.php?f=43.

2. Overview

We have already seen the camera board connected through the ribbon cable to (one of) the CSI port(s) on the Raspberry Pi itself.

Figure 2.
Overview of the libcamera system, explained below.



2.1. The Camera Module

Raspberry Pi already produces compatible camera boards, as listed in the introduction.

These are Bayer sensors, that is, they output so-called “raw” Bayer images which have not yet been processed into anything a user could recognise. These raw pixels are transmitted back to the Raspberry Pi in discrete image *frames*, using the MIPI CSI-2 protocol. This is the version of the CSI (Camera Serial Interface) defined by the MIPI (originally “Mobile Industry Processor Interface”) Alliance, an organisation dedicated to standardising interfaces between different components of mobile devices.

2.2. CSI Connector

Besides supplying power and clock signals, the ribbon cable performs two principal functions. It allows commands (typically register updates) to be sent to the image sensor using the $\mathrm{I^2\ C}$ interface. These values define the operating mode of the sensor, including the output image size and framerate, as well as commands to start and stop streaming images.

Secondly, the output images from the camera are transmitted through the cable along the CSI-2 interface back to the Raspberry Pi.

There are a number of different types of CSI-2 connectors. Raspberry Pi 4s and other older full-sized Raspberry Pis use the larger 15-pin connector. Compute Module 4s and Raspberry Pi 5s use smaller connectors accepting a 22-pin cable.

A more detailed specification of the modules and connector, including a link to the schematics, can be found at https://www.raspberrypi.com/documentation/accessories/camera.html.

2.3. On-Chip Hardware

The Raspberry Pi has the following hardware.

- An $\mathrm{I^2\ C}$ interface, for sending commands to the image sensor.

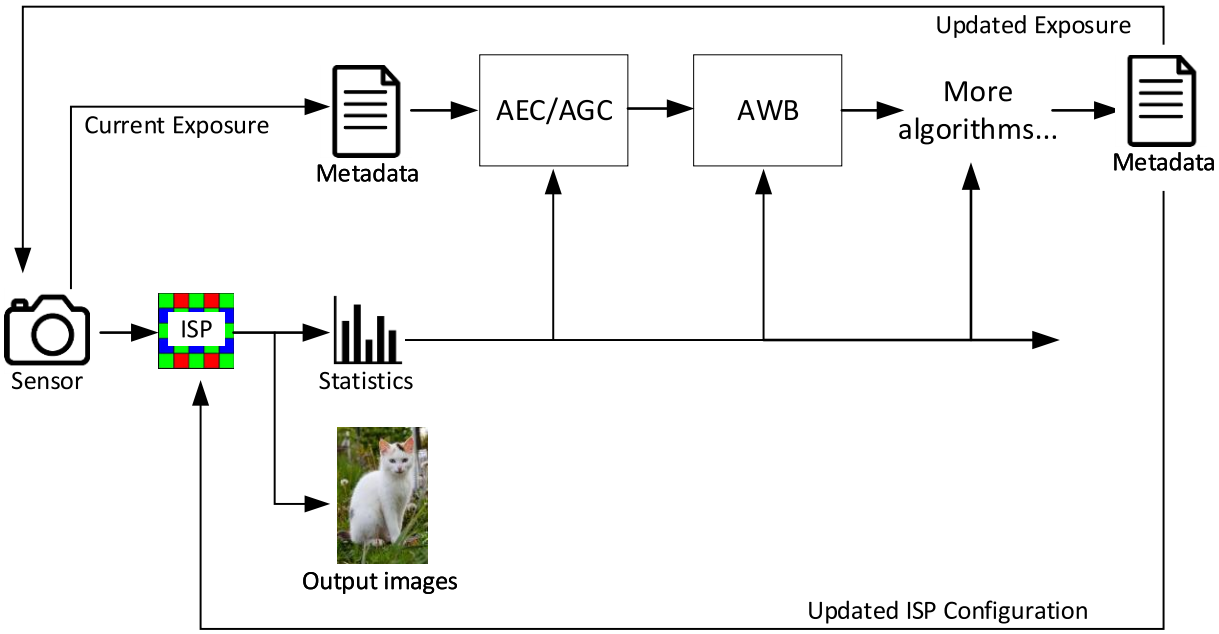
- A CSI-2 Receiver, for receiving image frames, in the form of pixel data, back from the sensor. The CSI-2 Receiver is on the main processor on Raspberry Pi 4 or older devices, where it is known as *Unicam*. On Raspberry Pi 5s, the CSI-2 receiver, along with some early "front end" parts of the ISP (Image Signal Processor) are on the RP1 I/O chip. Collectively, these entities are referred to as the *Camera Front End* (often abbreviated to CFE).
- And an ISP (Image Signal Processor) which converts the raw Bayer images from the sensor into something a user might recognise. On Pi 4 and earlier devices, the ISP also produces statistics from the pixel data, from which a number of software control algorithms will deduce appropriate parameters to be fed to the ISP in order to produce pleasing output images. On Pi 5, the statistics are produced by the CFE (Camera Front End), though the bulk of the processing (the *ISP Back End*) still happens on the main processor.

Interactions between these on-chip hardware devices is mediated though interrupts, responded to by the camera software stack.

2.4. Software and Control Algorithms

libcamera is responsible for handling and fulfilling application requests using Raspberry Pi's *pipeline handler*. In turn, the pipeline handler must ensure that the camera and ISP control algorithms are invoked at the correct moment, and must fetch up to date parameter values for each image frame.

Figure 3.
Control algorithms use image metadata and statistics to update their parameters.



The Raspberry Pi control algorithm framework associates a buffer of *metadata* with each image, containing exposure and other information about it. The algorithms each receive this metadata and statistics about the image which were calculated by the hardware ISP (or CFE on Pi 5). They run in turn to compute new and updated values which are deposited back into the metadata.

Finally, the pipeline handler reads values out of the metadata and uses them to update the camera's exposure settings and the ISP parameters for the next frame. There will be more details on this in chapter 4 (Control Algorithm Overview).

3. Driver Framework

There are a number of different types of driver involved in the camera system.

Firstly there are drivers for the various Raspberry Pi hardware devices, including Unicam (Pi 4 and earlier devices), the CFE (Camera Front End) on Pi 5, the ISP and so forth. These are all published in the `raspberrypi/linux` Github repository (https://github.com/raspberrypi/linux). Users and third parties should not need to change these drivers, but if you feel that you do and need help, please reach out on the Raspberry Pi support forums.

3.1. Camera Drivers

Those wishing to add support for new cameras will need to provide a standard Linux kernel driver for the new image sensor. Before embarking on this, we would recommend gaining some experience of Linux kernel drivers, and also of the Linux V4L2 framework.

Much the easiest route to creating a new camera driver is to start from an existing one, and to replace the calculations and register settings with ones that you have obtained from the sensor manufacturer. Both the imx477 and ov5647 would be suitable starting points.

You should endeavour to duplicate all the ioctls and V4L2 controls that you find in those drivers, as most of them will be mandatory for *libcamera* to work correctly.

Note that some sensors are able to return "embedded data" that lists the registers (including the exposure and gain registers) used for each frame. If your sensor is able to do this, then the imx477 driver is probably a better starting point. Otherwise we would recommend looking at the ov5647 driver first.

3.2. Device Tree

The camera driver needs to be listed as part of the Linux device tree so that it is correctly loaded. The details are largely beyond the scope of this document, however, we draw the reader's attention to the following resources.

1. In the `raspberrypi/linux` Github repository (https://github.com/raspberrypi/linux), please consult the `Documentation/devicetree/bindings/media/i2c` directory for example bindings, such as `imx219.txt` or `ov5647.txt`.
2. The existing overlays in the linux source tree at `arch/arm64/boot/dts/overlays` (use `arm` in place of `arm64` for 32-bit platforms). For example, look for `imx219-overlay.dts` or `ov5647-overlay.dts`.

3.3. The `CamHelper` Class

Unfortunately there is a certain amount of information that we need to know about cameras and camera modes but which the kernel driver framework (V4L2) does not supply. For this reason we provide the `CamHelper` class, along with derived class implementations for all of Raspberry Pi's supported sensors. These are all device-specific functions, but implemented in userspace rather than by extending the existing kernel framework. The principal functions of this class are:

1. To convert exposure times (in microseconds) to and from the device-specific representation (usually a number of lines of pixels) used by V4L2.
2. To convert analogue gain values to and from the device-specific "gain codes" that must be supplied to V4L2.
3. To allow parsing of embedded data buffers returned by a sensor.
4. To indicate how many frames of delay there are when updating the sensor's exposure time or analogue gain values.
5. To indicate how many frames are invalid and may need to be dropped on startup, or after changing the sensor mode and re-starting streaming.

Generally speaking it should be sufficient to copy one of our existing implementations and work from there. In particular note that

- Sensors that do not supply embedded data buffers should copy the OV5647 implementation (`cam_helper_ov5647.cpp`).
- Sensors that do supply embedded data buffers should copy the imx477 version (`cam_helper_imx477.cpp`). If the sensor in question follows the general SMIA pattern then this implementation, and the metadata parser, should in fact be quite close to what is required.

The `CamHelper` class methods are well documented in `cam_helper.hpp`, however, we do further draw the reader's attention to the `CameraMode` class, which contains the important parameters related to each different camera mode. The fields of the `CameraMode` are populated by the IPA, so they should automatically contain the correct values once the kernel driver is functioning properly. These fields documented in `controller/camera_mode.h`.

4. Control Algorithm Overview

4.1. Framework

Whilst Raspberry Pi provides a complete implementation of all the necessary ISP control algorithms, these actually form part of a larger framework. The principal components of this framework are described as follows.

- The `Controller` class is a relatively thin layer which manages all the control algorithms running under its *aegis*. It is responsible for loading the control algorithms at startup and invoking their methods when instructed to through *libcamera*.
- The `Algorithm` class is specialised to implement each individual control algorithm. The `Controller` invokes the methods of each algorithm as instructed.

The principal methods of the `Algorithm` class are

1. The `Prepare` method which is invoked just before the ISP is started. It is given information about the captured image (including shutter speed and amount of analogue gain applied by the sensor) and must decide the actual values to be programmed into the ISP.
2. The `Process` method which is invoked once statistics for the image frame are available. This method is given these statistics and is expected to initiate new computations which will determine the values written by the `Prepare` method when it runs again. The `Process` method runs just after the ISP has produced an output frame on Pi 4 and earlier devices, because it is the act of running the ISP that produces that statistics. On a Pi 5, however, statistics are collected by the Camera Front End, so `Process` runs just before `Prepare` , which programs the Back End ISP.

Other methods include the `Read` method, which loads the algorithm and may also initialise default values, and the `SwitchMode` method, which is invoked whenever a new `CameraMode` is selected for use.

In the following discussion, pathnames of specific files are given relative to the location of the Raspberry Pi specific code in the *libcamera* source tree, namely `libcamera/src/ipa/rpi`, unless we explicitly state otherwise.

4.2. Defining and Loading Algorithms

At startup, the `Controller` loads and initialises those algorithms that are listed in a JSON file. The idea is that a separate JSON file (the *camera tuning*) is loaded for each different kind of camera. Every algorithm is registered with the system using a static `RegisterAlgorithm` initialiser, making it easy to add more algorithms to the system.

The JSON file lists camera algorithms by name, conventionally given in a *<vendor>.<moniker>* format. The *moniker* can be thought of as the *type* of the algorithm, such as AWB (Automatic White Balance) or AGC/AEC (Automatic Gain/Exposure Control). For instance, Raspberry Pi's implementation of AWB is named `rpi.awb` . The implementation of each algorithm defines what parameters it has and in its `Read` method how it parses them from the file or sets default values.

If, for example, the mythical *Foo Corporation* were to produce their own implementation of AWB, it might be named `foo.awb` , and it could be listed directly in the JSON camera tuning file. Any algorithms no longer wanted could be deleted from the JSON file, or commented out by renaming them. For instance, changing `rpi.awb` to `x.rpi.awb` will cause the Raspberry Pi version to be silently ignored. We encourage developers to use a separate folder for any algorithms they implement, so while Raspberry Pi's algorithms all live in `controller/rpi`, Foo Corporation's algorithms might live in `controller/foo`.

Finally, we note that the order in which algorithms are listed in the JSON file is the order in which the controller will load and initialise them, and is also the order in which the `Prepare` and `Process` methods of the algorithms will run.

4.3. Standard Algorithms

The framework defines a number of abstract classes derived from `Algorithm` that provide standard interfaces to particular kinds of algorithms. For example, the `class AwbAlgorithm` (`controller/awb_algorithm.hpp`) requires the implementation of a `setMode` method which allows applications to set the AWB mode to, for example, `tungsten` , `daylight` , and so forth. Developers are encouraged to use these standard algorithm types when this is possible. Where it may be generally helpful, Raspberry Pi will also consider pull requests to add new features to the standard algorithms.

Furthermore, implementations of these standard algorithms are expected, by convention, to have a name which ends with the appropriate moniker for that type of algorithm. So an AWB algorithm would have a name always ending in `awb` , and AGC/AEC (Automatic Gain/Exposure Control) would always end in `agc` , facilitating the type of code shown below.

```
Code
Algorithm *algorithm = controller->GetAlgorithm("awb");
AwbAlgorithm *awb_algorithm = dynamic_cast<AwbAlgorithm *>(algorithm);
if (awb_algorithm) {
        awb_algorithm->setMode("sunny");
  }
```

The full set of standard monikers for our different types of algorithm is set out in the table below, though there is no reason in principle why developers should not invent new ones if they are appropriate for their own use cases. The algorithms themselves are explained in detail in the following chapter.

Moniker	Description
`alsc`	Automatic Lens Shading Correction algorithm.
`agc`	Automatic Gain/Exposure Control algorithm (AGC/AEC).
`awb`	Automatic White Balance algorithm (AWB).
`black_level`	Algorithm to set the correct black level for the sensor.
`cac`	Algorithm to mitigate colour fringing caused by lateral chromatic aberration. Pi 5 only.
`ccm`	Algorithm to calculate the correct Colour Correction Matrix (CCM).
`contrast`	Adaptive global contrast and gamma control algorithm.
`denoise`	Algorithm to jointly control spatial denoise, temporal denoise and colour denoise. Pi 5 only.
`dpc`	Algorithm to set the appropriate level of Defective Pixel Correction.
`geq`	Algorithm to set an appropriate level of Green Equalisation.
`lux`	Algorithm that estimates an approximate lux level for the scene.
`noise`	Algorithm to calculate the noise profile of the image given the current conditions.
`sdn`	Algorithm to set an appropriate strength for the Spatial Denoise block“. Pi 4 and earlier devices only.
`sharpen`	Algorithm to set appropriate sharpness parameters.
`af`	Algorithm to control a lens driver for Auto Focus.

4.4. Algorithm Communication

Algorithms communicate with the outside world, and with each other, through *image metadata*. Image metadata is an instance of the `Metadata` class (`controller/metadata.hpp`) which is created for each new image from the camera, and travels alongside the image through the entire processing pipeline; it is available to both the `Prepare` and `Process` methods. Code external to the algorithms themselves will look in the image metadata to find updated values to program into the ISP, and some algorithms may also look in the metadata to find the results calculated by other algorithms that they wish to use.

This information written into the image metadata is normally referred to as the algorithm's *status*, and is implemented by a class whose name ends in `Status` , normally found in the files `controller/*_status.h`. For example, the CCM algorithm (which calculates the colour correction matrix for the ISP) relies on the estimated colour temperature which can be found in the `AwbStatus` object (`controller/awb_status.h`) in the image metadata. In particular, items are stored in the image metadata indexed by a string key, and by convention this key is made up of two parts joined by a period (.), the first being the algorithm moniker (`awb`) and the second the word `status` .

Furthermore, we normally arrange that the status metadata contains any user-programmable settings that affect the result of the algorithm. This is so that an application can request that a setting be changed, and then sit back and monitor the image metadata to discover when it has actually happened. For example, we have seen the `setMode` method of the AWB algorithm. When the `AwbStatus` metadata shows the new mode name in its `mode` field, then we know the change has taken effect.

When replacing existing algorithms, care must be taken still to generate the correct status information in the image's metadata, as other (un-replaced) algorithms may still expect it. So any replacement AWB algorithm would still have to write out an `AwbStatus` object for the CCM algorithm (under the key `awb.status`), otherwise the CCM algorithm would need amending or replacing also. Custom algorithms are free to communicate using the image metadata too, only in this case the expectation is that any metadata

keys would begin with the organisation name. To wit, our friends at Foo Corporation might file their secret AWB parameters under `foo.awb.secret_parameters` .

Recall that the order in which `Prepare` and `Process` methods are run matches that of the JSON file, meaning that when one algorithm has a dependency on the status information of another then that determines their relative ordering within the file. So in the case of our example, we will have to list `rpi.ccm` **after** `rpi.awb` .

4.5. Performance Considerations

When writing new control algorithms, care must be taken not to delay the operation of the imaging pipeline any longer than necessary. This means that both the `Prepare` and `Process` methods for any algorithm must take considerably less than a millisecond, and that any significant amount of computation must be backgrounded using an asynchronous thread. Amongst the Raspberry Pi algorithms, examples of this can be found in the AWB (`controller/rpi/awb.cpp`) and ALSC (Automatic Lens Shading Correction - `controller/rpi/alsc.cpp`) algorithms.

Although the details are of course up to developers, the general pattern is for the `Process` method to re-start the computation (in the asynchronous thread) if the previous computation has finished. Often it makes sense to wait a few frames before re-starting it as the results are not usually so critical that we need them immediately, and we can thereby save on CPU load too.

The `Prepare` method normally has a notion of *target values*, and it filters the values that it outputs slowly towards these target values with every frame (for example, using an IIR filter). It watches for any in-progress asynchronous computation to finish though it does not block" waiting for it. If unfinished, it will simply leave the target values unchanged and try again on the next frame; if finished, this merely causes an update to its target values, which it will immediately start filtering towards.

4.6. Example Camera Tuning File

4.6.1. Pi 4 and Earlier Devices

For earlier Raspberry Pi devices, the "target" field should read "bcm2835".

Code

```
{
    "version": 2.0,
    "target": "bcm2835",
    "algorithms": [
        {
            "rpi.black_level":
            {
                "black_level": 4096
            }
        },
        {
            "rpi.awb":
            {
                "use_derivatives": 0,
                "bayes": 0
            }
        },
        {
            "rpi.agc":
            {
                "metering_modes":
                {
                    "centre-weighted":
                    {
                        "weights": [ 4, 4, 4, 2, 2, 2, 2, 1, 1, 1, 1, 0, 0, 0, 0 ]
                    }
                },
                "exposure_modes":
                {
                    "normal":
```

```
                {
                    "shutter": [ 100, 15000, 30000, 60000, 120000 ],
                    "gain": [ 1.0, 2.0, 3.0, 4.0, 6.0 ]
                },
                "short":
                {
                    "shutter": [ 100, 5000, 10000, 20000, 30000 ],
                    "gain": [ 1.0, 2.0, 4.0, 6.0, 8.0 ]
                },
                "long":
                {
                    "shutter": [ 1000, 30000, 60000, 90000, 120000 ],
                    "gain": [ 1.0, 2.0, 4.0, 6.0, 12.0 ]
                }
            },
            "constraint_modes":
            {
                "normal": [
                    {
                        "bound": "LOWER",
                        "q_lo": 0.98,
                        "q_hi": 1.0,
                        "y_target":
                        [
                            0, 0.4,
                            1000, 0.4
                        ]
                    }
                ]
            },
            "y_target":
            [
                0, 0.16,
                1000, 0.165,
                10000, 0.17
            ]
        }
    },
    {
        "rpi.ccm":
        {
            "ccms": [
                {
                    "ct": 4000,
                    "ccm":
                    [
                        2.0, -1.0, 0.0,
                        -0.5, 2.0, -0.5,
                        0, -1.0, 2.0
                    ]
                }
            ]
        }
    },
    {
        "rpi.contrast":
        {
            "ce_enable": 0,
            "gamma_curve":
            [
                0, 0,
                1024, 5040,
                2048, 9338,
                3072, 12356,
                4096, 15312,
```

```
                    5120, 18051,
                    6144, 20790,
                    7168, 23193,
                    8192, 25744,
                    9216, 27942,
                    10240, 30035,
                    11264, 32005,
                    12288, 33975,
                    13312, 35815,
                    14336, 37600,
                    15360, 39168,
                    16384, 40642,
                    18432, 43379,
                    20480, 45749,
                    22528, 47753,
                    24576, 49621,
                    26624, 51253,
                    28672, 52698,
                    30720, 53796,
                    32768, 54876,
                    36864, 57012,
                    40960, 58656,
                    45056, 59954,
                    49152, 61183,
                    53248, 62355,
                    57344, 63419,
                    61440, 64476,
                    65535, 65535
                ]
            }
        }
    ]
}
```

4.6.2. Pi 5 Devices

For devices based on Raspberry Pi's new ISP, the "target" field must read "pisp". The format of the file is very similar, though in places it differs on account of the different capabilities of the underlying hardware.

Code

```
{
    "version": 2.0,
    "target": "pisp",
    "algorithms": [
        {
            "rpi.black_level":
            {
                "black_level": 4096
            }
        },
        {
            "rpi.awb":
            {
                "use_derivatives": 0,
                "bayes": 0
            }
        },
        {
            "rpi.agc":
            {
                "metering_modes":
                {
                    "centre-weighted":
                    {
```

```
                    "weights":
                    [
                        0, 0, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 0, 0,
                        0, 1, 1, 1, 1, 1, 2, 2, 2, 1, 1, 1, 1, 1, 0,
                        1, 1, 1, 1, 2, 2, 2, 2, 2, 2, 2, 1, 1, 1, 1,
                        1, 1, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 1, 1,
                        1, 1, 2, 2, 2, 2, 3, 3, 3, 2, 2, 2, 2, 1, 1,
                        1, 1, 2, 2, 2, 3, 3, 3, 3, 3, 2, 2, 2, 1, 1,
                        1, 1, 2, 2, 3, 3, 3, 4, 3, 3, 3, 2, 2, 1, 1,
                        1, 1, 2, 2, 3, 3, 4, 4, 4, 3, 3, 2, 2, 1, 1,
                        1, 1, 2, 2, 3, 3, 3, 4, 3, 3, 3, 2, 2, 1, 1,
                        1, 1, 2, 2, 2, 3, 3, 3, 3, 3, 2, 2, 2, 1, 1,
                        1, 1, 2, 2, 2, 2, 3, 3, 3, 2, 2, 2, 2, 1, 1,
                        1, 1, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 1, 1,
                        1, 1, 1, 1, 2, 2, 2, 2, 2, 2, 2, 1, 1, 1, 1,
                        0, 1, 1, 1, 1, 1, 2, 2, 2, 1, 1, 1, 1, 1, 0,
                        0, 0, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 0, 0
                    ]
                }
            },
            "exposure_modes":
            {
                "normal":
                {
                    "shutter": [ 100, 15000, 30000, 60000, 120000 ],
                    "gain": [ 1.0, 2.0, 3.0, 4.0, 6.0 ]
                }
            },
            "constraint_modes":
            {
                "normal": [
                    {
                        "bound": "LOWER",
                        "q_lo": 0.98,
                        "q_hi": 1.0,
                        "y_target":
                        [
                            0, 0.4,
                            1000, 0.4
                        ]
                    }
                ]
            },
            "y_target":
            [
                0, 0.16,
                1000, 0.165,
                10000, 0.17
            ]
        }
    },
    {
        "rpi.ccm":
        {
            "ccms": [
                {
                    "ct": 4000,
                    "ccm":
                    [
                        2.0, -1.0, 0.0,
                        -0.5, 2.0, -0.5,
                        0, -1.0, 2.0
                    ]
                }
            ]
```

```
            }
        },
        {
            "rpi.contrast":
            {
                "ce_enable": 0,
                "gamma_curve":
                [
                    0, 0,
                    1024, 5040,
                    2048, 9338,
                    3072, 12356,
                    4096, 15312,
                    5120, 18051,
                    6144, 20790,
                    7168, 23193,
                    8192, 25744,
                    9216, 27942,
                    10240, 30035,
                    11264, 32005,
                    12288, 33975,
                    13312, 35815,
                    14336, 37600,
                    15360, 39168,
                    16384, 40642,
                    18432, 43379,
                    20480, 45749,
                    22528, 47753,
                    24576, 49621,
                    26624, 51253,
                    28672, 52698,
                    30720, 53796,
                    32768, 54876,
                    36864, 57012,
                    40960, 58656,
                    45056, 59954,
                    49152, 61183,
                    53248, 62355,
                    57344, 63419,
                    61440, 64476,
                    65535, 65535
                ]
            }
        }
    ]
}
```

The contents of the above JSON file are not particularly important for the moment, however we note the use of named Raspberry Pi (`rpi`) algorithms. This example file is deliberately very simple, but contains enough information to give quite recognisable images from almost any sensor.

4.7. Directory Structure

A summary of the folders and files provided as part of the Raspberry Pi framework is given below. The paths listed can be assume to lie under *libcamera*'s `src/ipa/rpi` folder unless stated otherwise.

Path	Description
`cam_helper/cam_helper*`	Implementation of the `CamHelper` class for different sensors.
`cam_helper/md_parser_smia.cpp`	Functionality for parsing camera embedded data (see chapter 3).
`controller*_status.h`	Definitions of the status objects placed into the image metadata by control algorithms.

Path	Description
controller*_algorithm.hpp	Interfaces for standard algorithms, such as AWB, AGC/AEC etc.
controller/metadata.hpp	Class representing image metadata.
controller/histogram.cpp	A simple histogram class that may be useful to control algorithm implementations.
controller/pwl.cpp	A simple implementation of piecewise linear functions that may be useful to control algorithms. Many of Raspberry Pi's own control algorithms depend on these.
controller/rpi	Folder containing all of Raspberry Pi's algorithm implementations.
vc4/data	Folder containing existing camera tuning JSON files for Pi 4 and earlier devices.
pisp/data	Folder containing existing camera tuning JSON files for Pi 5 devices.

5. Raspberry Pi Control Algorithms

This section documents Raspberry Pi's implementation of the ISP and camera control algorithms. For each algorithm we give an overview and if necessary explain any concepts underlying the algorithm. We list all of the tuning parameters for each algorithm including what they mean and their default values, and explicitly describe any metadata dependencies. Note further that when any of these algorithms runs, the image metadata already contains the `DeviceStatus` metadata (`controller/device_status.h`), giving the shutter speed and analogue gain of the current image.

The algorithms are listed in the order in which they would normally appear in the JSON file as this is probably the most natural order in which to understand them.

Many of the parameters listed here can be left at default values, however, others may appear distinctly esoteric. Users are reminded that a Camera Tuning Tool (CTT) is provided (the subject of chapter 6) which facilitates the automatic generation of JSON tuning files with all these difficult values correctly calibrated.

5.1. Black Level

5.1.1. Name

`rpi.black_level`

5.1.2. Overview

The pixels output by the camera normally include a "pedestal" value or *black level*. Pixels at or below this level should be considered black. For a camera outputting 10-bit pixel values (in the range 0 to 1023) a typical black level might be 64, but the correct value would be found in the sensor manufacturer's datasheet. The number entered into the tuning file here is agnostic of the sensor's native bit depth and always expects a value from a 16-bit dynamic range. So in this case, we would supply the value 4096 $(= 64/1024 \times 65536)$.

Name	Default	Description
`black_level`	4096	Specifies a default black level value to use for all three RGB channels, unless over-ridden by subsequent entries.
`black_level_r`	*Optional*	If specified, over-rides the default `black_level` for the red channel.
`black_level_g`	*Optional*	If specified, over-rides the default `black_level` for the green channel.
`black_level_b`	*Optional*	If specified, over-rides the default `black_level` for the blue channel.

5.1.3. Metadata Dependencies

The Black Level control algorithm has no dependency on any other image metadata.

The Black Level control algorithm writes a `BlackLevelStatus` object (`controller/black_level_status.h`) into the image metadata. This gives a (16-bit) black level value for each of the three RGB channels.

5.2. Defective Pixel Correction (DPC)

5.2.1. Name

`rpi.dpc`

5.2.2. Overview

Camera images may contain a number of “stuck” or sometimes “weak” pixels. Pixels in the former category are stuck at a particular value, sometimes completely black or sometimes completely white, whereas the latter kind (weak pixels) vary atypically according to the incoming light and can be brighter or darker than their neighbours. In both cases this can lead to undesirable image artifacts.

Figure 4.
Magnified image showing a bright defective pixel - its effect spreads slightly because of the filtering operations performed using it.



The ISP hardware has the ability to filter some of these out adaptively, without knowing in advance where they are. We expose a simple control for this, namely *no correction*, *normal correction* or *strong correction*.

5.2.3. Parameters

Name	Default	Description
`strength`	1	Strength of the correction. 0 = no correction, 1 = “normal” correction and 2 = “strong” correction.

5.2.4. Metadata Dependencies

The DPC control algorithm has no dependency on any other image metadata.

The DPC control algorithm writes a `DpcStatus` object (`controller/dpc_status.h`) into the image metadata.

5.3. Lux

5.3.1. Name

`rpi.lux`

5.3.2. Overview

The Lux algorithm does not actually generate any parameters that are subsequently programmed into the ISP. Instead, it merely provides a single location where an estimate of the lux level (amount of light) of the current image can be generated. It writes this value into the image metadata where downstream control algorithms can find it and use it.

It is set up with the exposure/gain parameters and known lux level of a reference image, which must be calibrated during the camera tuning process (most straightforwardly with the CTT). It generates its lux estimate for the current image by a simple ratio calculation comparing current exposure/gain and statistics values to the ones recorded for the reference image.

5.3.3. Parameters

Name	Default	Description
`reference_shutter_speed`	*Required*	Shutter speed (exposure time) of the reference image, in microseconds, e.g. `33000` .
`reference_gain`	*Required*	Analogue gain of the reference image, e.g. `1.5` .
`reference_Y`	*Required*	Measured Y (luminance) value of the reference image. This should normally be obtained through the CTT.
`reference_lux`	*Required*	Estimated lux level of the reference image scene. Normally supplied by the CTT.

5.3.4. Metadata Dependencies

The Lux control algorithm requires the presence of the `DeviceStatus` metadata (which should automatically always be available).

The Lux control algorithm writes a `LuxStatus` object (`controller/lux_status.h`) into the image metadata.

5.4. Noise

5.4.1. Name

`rpi.noise`

5.4.2. Overview

The Noise algorithm also does not calculate any new parameters for the ISP, it merely calculates the *noise profile* for the current image. It does this by being calibrated with the noise profile of a reference image and then, for each image passing through the pipeline, adjusts the noise profile according to the analogue gain of the current image and also information about the *camera mode*. The camera mode is significant because most cameras are able to average pixels together, outputting a less noisy image than the native full resolution image. These values would normally be calculated by the CTT (Camera Tuning Tool).

5.4.3. Parameters

Name	Default	Description
`reference_constant`	*Required*	Constant offset of the noise profile. In practice, it can be taken always to have the value zero.
`reference_slope`	*Required*	Slope of the noise profile, plotted against the square root of the pixel value. Supplied by the CTT, a typical value might be somewhere around `3.0` .

5.4.4. Metadata Dependencies

The Noise control algorithm requires the presence of the `DeviceStatus` metadata (which should automatically always be available).

The Noise control algorithm writes a `NoiseStatus` object (`controller/noise_status.h`) into the image metadata.

5.5. GEQ (Green Equalisation)

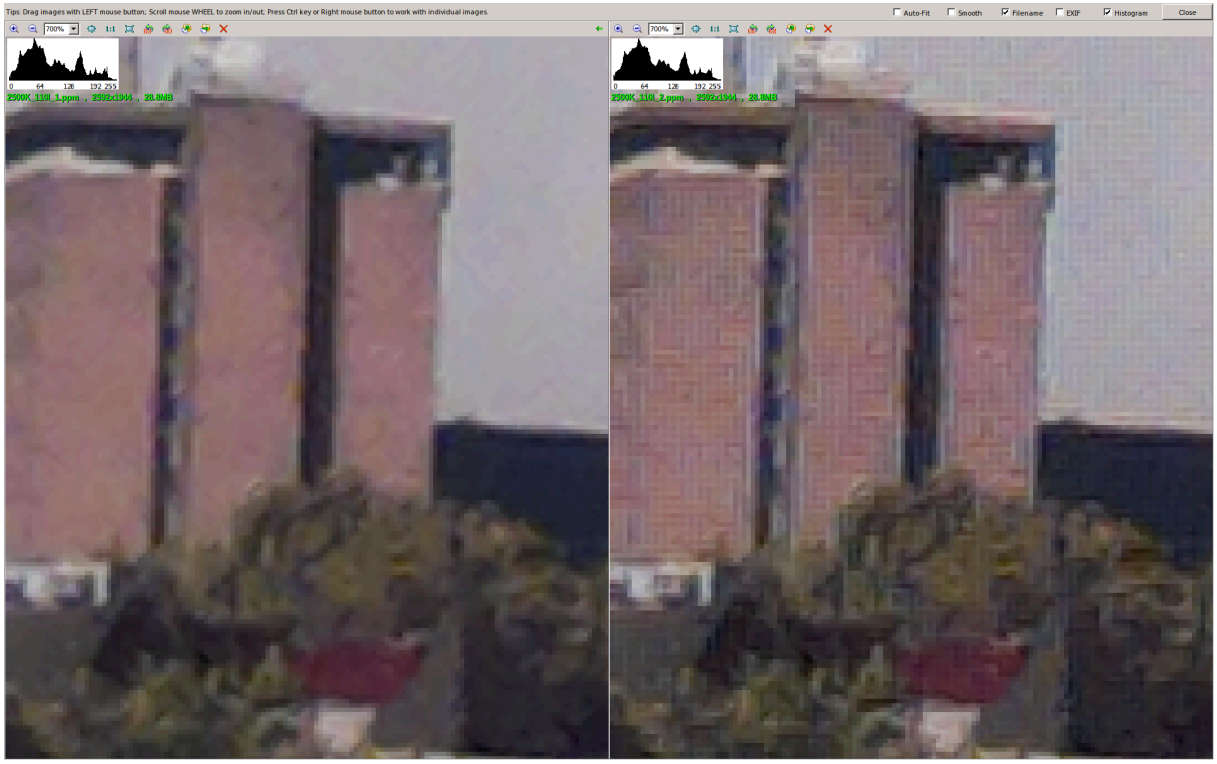
5.5.1. Name

`rpi.geq`

5.5.2. Overview

Standard Bayer cameras produce raw images with two slightly different kinds of green pixels, those on the red rows, and those on the blue rows. Under some circumstances these can respond slightly differently even when exposed to a solid and completely uniform colour. Downstream processing (especially *demosiacking*) can have a tendency to pick up and accentuate these differences, giving rise to what are frequently described as *maze artifacts*.

Figure 5.
Highly magnified image showing maze artifacts (right), and correction (left).



Green Equalisation is the process of correcting green imbalance, shown in the image above. Most simply, it is applying a low pass filter to the green pixels, but it needs to be given some expectation of the size of the imbalance so as not to blur the image more than is necessary. The Green Equalisation parameters would normally be provided by the CTT.

5.5.3. Parameters

Name	Default	Description
`offset`	0	Constant offset of the green imbalance level.
`slope`	0.0	Slope of the green imbalance level (compared to pixel value). Must be < 1.0.
`strength`	*Optional*	If specified, a piecewise linear function that applies a gain to the green imbalance level based on estimated lux level. This allows even stronger green imbalance correction in very low light situations.

5.5.4. Metadata Dependencies

The Green Equalisation algorithm requires the presence of the `DeviceStatus` metadata (which should automatically always be available). It also requires the presence of `LuxStatus` metadata from the Lux control algorithm.

The Green Equalisation algorithm writes a `GeqStatus` object (`controller/geq_status.h`) into the image metadata.

5.6. SDN (Spatial Denoise)

5.6.1. Name

`rpi.sdn`

Note: this algorithm applies only to the Pi 4 and earlier devices. For Pi 5, please refer to the `rpi.denoise` algorithm.

5.6.2. Overview

The SDN (Spatial Denoise) control algorithm generates parameters for controlling the ISP's spatial denoise function. It derives a new noise profile from the one calculated (and placed in the image metadata) by the Noise control algorithm, and this is ultimately passed to the ISP. It also determines the "strength" of the denoise function, meaning how much of the final pixel is taken from the denoised pixel, the remainder being made up by the un-denoised pixel. This turns out to be helpful because denoising an image completely can leave it looking a little artificial or "plasticky"; leaving some noise in the final image normally proves to be aesthetically beneficial.

The SDN control algorithm parameters can normally just be left at their default values.

Figure 6.
Same image as figure 1, highly magnified, showing "noise speckles".



Colour channel denoising is another operation controlled by the SDN algorithm. Colour denoising can operate in two modes - a high quality mode, or a fast mode. The former should be used for still image captures where frame time is not critical, while the latter should be used for viewfinder and video recoding where processing time is more important to avoid dropping frames. No colour denoising is enabled by default.

5.6.3. External API

The `Sdn` class is derived from the `DenoiseAlgorithm` class. As such it defines the following publicly accessible methods.

DenoiseAlgorithm class API
`setMode(DenoiseMode mode)` Sets the operating mode for SDN. `DenoiseMode` may be one of the following enum values { `Off` , `ColourOff` , `ColourFast` , `ColourHighQuality` }.

An example use of the DenoiseAlgorithm method:

```
Code
Algorithm *algorithm = controller->GetAlgorithm("sdn");
DenoiseAlgorithm *denoise_algorithm = dynamic_cast<AgcAlgorithm *>(algorithm);
if (denoise_algorithm) {
    denoise_algorithm->setMode(DenoiseMode::ColourHighQuality)
}
```

5.6.4. Parameters

Name	Default	Description
`deviation`	3.2	Scaling applied to the Noise control algorithm's noise profile. Increasing this means wider pixel differences will be denoised (producing a blurrier image) and reducing it results in only smaller pixel differences being denoised. Making this too small will produce an image with little "noise speckles" in it.
`strength`	0.75	Proportion of the final output pixel made up by the denoised pixel value, the remainder (0.25 in the default case) coming from the original (noisy) pixel.

5.6.5. Metadata Dependencies

The SDN control algorithm requires the presence of the `NoiseStatus` metadata from the Noise control algorithm.

The SDN control algorithm writes a `DenoiseStatus` object (`controller/denoise_status.h`) into the image metadata.

5.7. Denoise (Combined Denoise Algorithms)

5.7.1. Name

`rpi.denoise`

Note: this algorithm applies only to the Pi 5. For earlier devices, please refer to the `rpi.sdn` algorithm.

5.7.2. Overview

The Denoise control algorithm generates parameters jointly for the ISP's three denoise functions: SDN (Spatial Denoise), TDN (Temporal Denoise) and CDN (Colour Denoise). It derives a new noise profile from the one calculated (and placed in the image metadata) by the Noise control algorithm, and writes SDN, TDN and CDN control parameters into the image metadata in order to be programmed into the hardware pipeline.

SDN performs spatial denoise, which involves averaging together pixels in a small neighbourhood. Where pixels look "close enough" it will average them.

TDN (Temporal Denoise) averages a pixel in the frame that has arrived most recently with the previous output of the TDN block (that is, the temporally denoised version of the previous frame). Where they look similar, they are averaged, and where they do not, the most recent pixel is taken. TDN denoises more effectively once it has seen "several frames".

Finally CDN (Colour Denoise) is another spatial filter but works specifically on the colour, rather than brightness, signal. Here the scale of the noise tends to cover a larger area, but we can also denoise more aggressively because the colour signal contributes less to image detail.

On Pi 5, all these algorithms operate directly in hardware at the full pixel rate of the ISP.

5.7.3. External API

The `Denoise` class is derived from `DenoiseAlgorithm`, defining the following publicly accessible methods.

DenoiseAlgorithm class API
`setMode(DenoiseMode mode)` Sets the operating mode for SDN. `DenoiseMode` may be one of the following enum values { `Off`, `ColourOff`, `ColourFast`, `ColourHighQuality` }. Pi 5 does not distinguish between “fast” and “high quality” colour denoise as the two are the same on this platform.
`setConfig(std::string const &name)` Sets the current denoise configuration. Different imaging modes may use slightly different denoise configurations such as `normal` for normal use-cases, `hdr` for HDR imaging and `night` for night mode imaging.

An example use of the DenoiseAlgorithm:

```
Code
Algorithm *algorithm = controller->GetAlgorithm("denoise");
DenoiseAlgorithm *denoise_algorithm = dynamic_cast<DenoiseAlgorithm *>(algorithm);
if (denoise_algorithm) {
    denoise_algorithm->setConfig("normal")
}
```

5.7.4. Parameters

Name	Default	Description
`sdn`		The following parameters should be in a `sdn` sub-block“.
`deviation`	3.2	Scaling applied to the Noise control algorithm’s noise profile. Increasing this means wider pixel differences will be denoised (producing a blurrier image) and reducing it results in only smaller pixel differences being denoised. Making this too small will produce an image with little “noise speckles” in it.
`strength`	0.75	Proportion of the final output pixel made up by the denoised pixel value, the remainder (0.25 in the default case) coming from the original (noisy) pixel.
`deviation2`	`deviation`	Deviation value used for a pixel that has not been temporally denoised. When temporal denoise is running, it can be helpful to set `deviation` lower because, as temporal denoise is running, pixels tend to be less noisy. But where TDN has decided not to denoise a pixel, a higher value here can be beneficial.
`deviation_no_tdn`	`deviation`	The SDN deviation value to use when the camera starts or switches mode. When this happens, TDN has to be reset and is therefore having a minimal effect for the first few frames. As frames pass through the system, the value passed to the hardware ramps down from this value to the normal `deviation` value.
`strength_no_tdn`	`strength`	The SDN strength value ramps from this value to the normal `strength` value as TDN operates on frames.
`backoff`	0.75	Controls the rate at which SDN backs off from the “no TDN” to the “normal” values with every frame. Larger values imply slower back-off.
`cdn`		The following parameters should be in a `cdn` sub-block“.
`deviation`	120	The colour denoise deviation parameter. Increasing it will cause more cause more colour denoise to happen.
`strength`	0.2	Colour denoise filter strength. A larger value implies more colour denoise.
`tdn`		The following parameters should be in a `tdn` sub-block“.
`deviation`	0.5	The temporal denoise deviation parameter. Increasing it will cause more cause more temporal denoise to happen.

Name	Default	Description
`threshold`	0.75	Temporal denoise threshold, below which pixels are considered not to have been temporally denoised. Increasing this value will cause more pixels to be treated by SDN as though temporal denoise did not change them, so that the `deviation2` and `strength2` parameters are applied to them.

5.7.5. Metadata Dependencies

The denoise control algorithm requires the presence of the `NoiseStatus` metadata from the Noise control algorithm.

The denoise control algorithm writes `SdnStatus` , `TdnStatus` and `CdnStatus` objects (`controller/denoise_status.h`) into the image metadata, giving the control parameters for each of the 3 denoise functions.

5.8. AWB (Automatic White Balance)

5.8.1. Name

`rpi.awb`

5.8.2. Overview

Readers are no doubt aware how, when you look at a particular (non-changing) object, its colour always appears the same, irrespective of whether you look at it outside in broad daylight or indoors under some kind of artificial light. If it looks *orange* outdoors, then it looks *orange* indoors too. Yet if you measured the colour directly using a spectrometer, you'd discover that *outdoors* the object is much *bluer* and *indoors* it is much *redder*.

The Human Visual System has adapted amazingly over time to hide these differences from us - a phenomenon known as *Colour Constancy*. Unfortunately no such magic occurs when viewing a digital image so we have to make this correction ourselves. This is the job of the AWB (Automatic White Balance) Algorithm. We describe the algorithm below, though readers should note that the settings and calibrations should normally be produced by the Camera Tuning Tool.

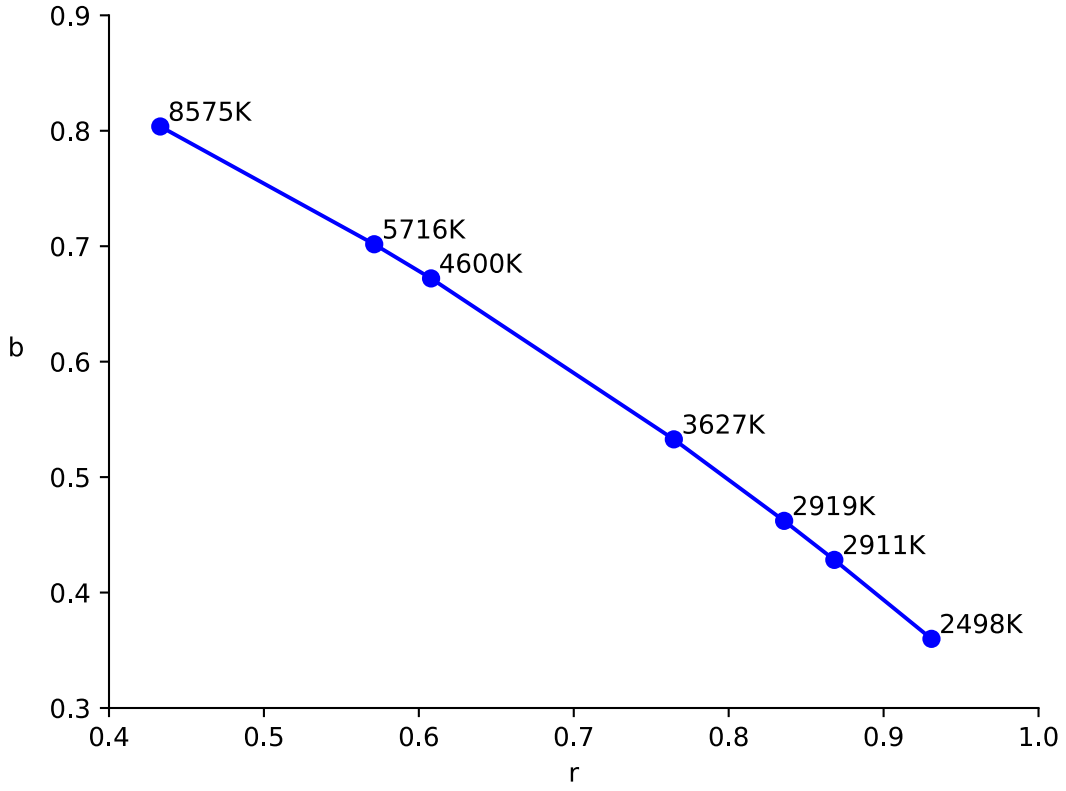
Raspberry Pi takes a simple Bayesian approach to the problem. We need to define a *feasible* set of illuminants and then, for any given image, choose that illuminant from the set which gives us the *most likely* resulting colours. The mathematically inclined will recognise this as a *maximum likelihood estimation* problem. Finally we apply the correction to the image according to the selected illuminant.

5.8.3. The Feasible Set

The illuminants that we encounter in the real world have historically tended to lie on the so-called *Planckian Locus*, as the light is produced by some kind of heat-radiating body (be it the sun, or a tungsten filament). In recent years, with the advent of fluorescent, LED and other light sources, this is less true, however, with some slight loosening of our constraints it can still be considered an adequate approximation.

The feasible set is constructed by measuring the response of the sensor to a known grey target. In practice this means capturing an image of something like a Macbeth Chart (with objectively grey squares) under each of our possible illuminants, and measuring the average RGB value of the grey squares. For each RGB triple we calculate normalised colour values $r = R/G$ and $b = B/G$, and plot these on a graph called the *Colour Temperature Curve* (or *CT Curve*) for the camera. Each point on the curve can be associated with the colour temperature of the illuminant that provided it.

Figure 7.
CT Curve showing colour temperature associated with each point.



Note how, in this example, colour temperatures get closer at the left hand (daylight) end of the curve, and spread out at the right hand (indoor) end. In particular, the points for colour temperatures 2911K and 2919K are quite distinct even though the colour temperatures are virtually identical. This is most likely because they were produced by different kinds of illuminant - perhaps one was a fluorescent tube and the other an incandescent bulb.

Finally, whilst the (r, b) coordinate gives the response of the camera to light of that colour temperature, the red and blue gains that need to be applied to the image in order to white balance it correctly (that is, make actual grey have an equal response on all RGB channels) are

$$gain_r = \frac{1}{r}$$

and

$$gain_b = \frac{1}{b}$$

5.8.4. Bayes Theorem

In our case we are wishing to find the illuminant I with the highest probability given the observed pixel data D, written as $P(I \mid D)$. According to Bayes' Theorem we have

$$P(I \mid D) = \frac{P(D \mid I) \cdot P(I)}{P(D)}$$

where

$P(I)$ is the prior probability of the illuminant I, $P(D)$ is the prior probability of the observed pixel data, $P(D \mid I)$ is the posterior probability of the observed pixel data D given the illuminant I, and finally $P(I \mid D)$ is the posterior probability of the illuminant I given the observed pixel data D.

5.8.5. The Priors

Firstly, in our case we take $P(D)$ as constant, therefore the problem reduces to finding

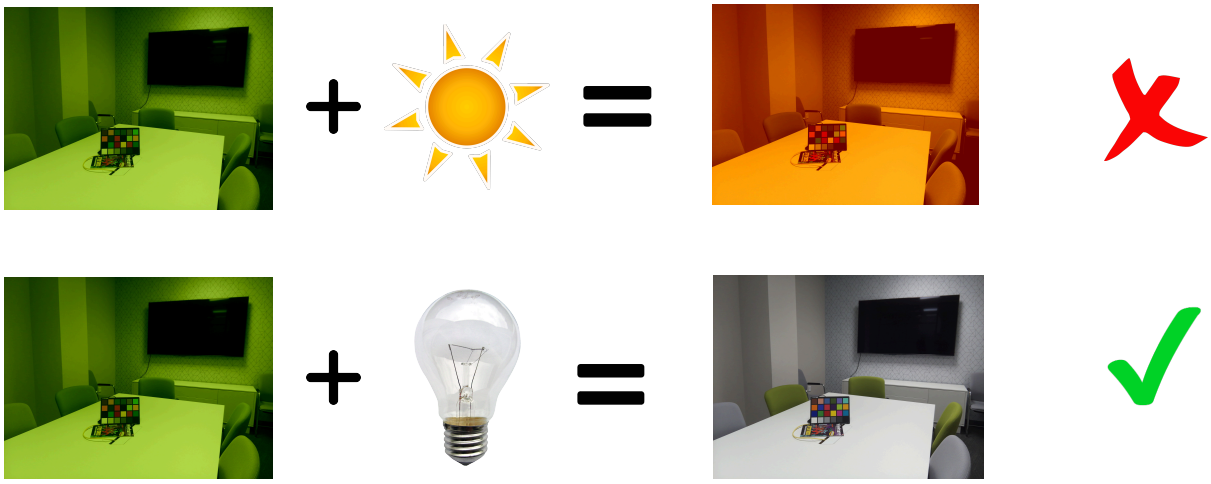
It is up to us how we wish to define $P(I)$, and indeed this is one of the principal means of tuning the AWB algorithm. Typically, we might choose larger (more likely) values for daytime illuminants when the ambient conditions (as measured by the Lux control algorithm) suggest the scene is very bright. Accordingly, the indoor illuminants would be dispreferred under these conditions. In medium to low light situations we would most likely reverse these biases, favouring the indoor illuminants instead.

5.8.6. Calculating the Posteriors

$P(D \mid I)$ is not difficult to calculate providing we have some kind of probabilistic model of the colours we expect to see in an output image. In our case we regard certain colours as more likely than others. In particular the most likely colour is generally grey or near-grey (the correspondence with simple “Grey World” algorithms should be clear), but we could consider others too - blue sky, green grass and skin tones might be the most likely candidates. (Note, however, that Raspberry Pi’s implementation only looks for near-grey colours at this time.) The procedure is then to take the raw camera image, apply the red and blue gains for a particular illuminant, and then assess how plausible the result looks.

We illustrate this below with an indoor image taken under incandescent lighting; note that the raw camera image (as is often the case) looks rather green.

Figure 8.
AWB example - the incandescent illuminant gives a more plausible result than the daylight one.



We also need to specify how we calculate the probability of any pixel’s colour. Given the pixel’s RGB colour values we calculate, as before, $r = R/G$ and $b = B/G$. For a given illuminant I (and hence known $gain_r$ and $gain_b$ values), we define colour difference statistics

$$\Delta r = gain_r \cdot r - 1$$

$$\Delta b = gain_b \cdot b - 1$$

Observe how, for the illuminant which turns the pixel grey we have that $\Delta r = \Delta b = 0$.

Now we assume that a 2-dimensional Gaussian distribution gives a reasonable approximation to the PDF (probability density function) of our variables Δr and Δb, meaning that for a single pixel we have

$$P(D \mid I) \propto e^{-\left(\frac{\Delta r^2}{\sigma_r^2} + \frac{\Delta b^2}{\sigma_b^2}\right)}$$

where σ_r and σ_b are the standard deviations of the distribution, and we use $\propto$ in place of equality because we have elided the (unimportant) fixed constant (that makes the PDF integrate to unity).

Now, we do not actually work directly on pixels as this would be far too slow. Instead the image is divided into $16 \times 12 = 192$ regions for which the hardware ISP calculates R, G and B averages; in many ways it can be thought of as if we are operating on an image only 16×12 pixels in size. So extending our probability calculation to the entire image, and assuming that all the regions are independent (an assumption we shall return to later), we obtain the following expression for the posterior probability of the illuminant:

$$P(I \mid D) \propto P(I) \prod_{regions} e^{-\left(\frac{\Delta r^2}{\sigma_r^2} + \frac{\Delta b^2}{\sigma_b^2}\right)}$$

5.8.7. Log Likelihood

At this point it should be clear that working in terms of *log likelihood* is going to be more convenient. The large product turns simply into a sum, and as we are defining $P(I)$ ourselves anyway we may as well just define its log likelihood instead. We use L in place of P to denote log likelihoods.

Further, we take $\sigma_r = \sigma_b$ as it actually makes little difference to the end results, and then can drop this term entirely as we can scale our $L(I)$ function (the prior log likelihood of the illuminant) to compensate. Finally we are left with finding

5.8.8. Implementation

Some notes on the Raspberry Pi implementation of the algorithm described above. It can be found in `controller/rpi/awb.cpp`.

1. The prior for the illuminant depends only on estimated lux level. Normally this involves interpolating between the two priors either side of the current lux estimate.
2. The CT Curve definition in the tuning file makes 2 piecewise linear functions, giving us r and b in terms of the colour temperature T. Thus $r = r(T)$ and $b = b(T)$.
3. The algorithm starts with a coarse search taking relatively large steps down the CT curve.
4. The "white point" the algorithm is looking for, in $(\Delta r, \Delta b)$ space is normally given by $(0, 0)$. However the configuration parameters (below) do allow this to be moved in case, for example, a slightly warmer look is desired.
5. Subsequently there is a fine search around the area found. The fine search is allowed to wander transversely off the CT Curve (within limits) to better accommodate lamps that do not lie perfectly on the curve (often including fluorescent lights and so forth).
6. As mentioned previously, the algorithm only expects neutral colours at this time.
7. The squared terms in the log likelihood expression get clamped; this reflects the fact that there genuinely are non-neutral colours in the world, and once something looks non-neutral it is counterproductive to penalise it further just because it appears *very* non-neutral.
8. For when no camera calibration has yet been done, our implementation includes a very plain Grey World algorithm that gives adequate results in many circumstances. When the Grey World algorithm is used, the estimated colour temperature produced by the AWB algorithm is not meaningful.

5.8.9. External API

The `Awb` class is derived from `AwbAlgorithm` . As such it defines the following publicly accessible methods.

AwbAlgorithm class API
`GetConvergenceFrames()` This returns the number of frames the AWB algorithm recommends be dropped, while it converges, when the camera and the AWB algorithm are started from scratch.
`SetMode(std::string const &mode_name)` This method takes the name of the desired AWB mode (a string such as `tungsten` or `daylight`). The effect is to limit the range of the coarse search to the range listed alongside the mode (which must be preset) in the JSON file.
`SetManualGains(double manual_r, double manual_b)` Switches to using the provided manual red and blue gains. Set both values to zero to switch back to auto mode.

An example use of AWB methods:

```
Code
Algorithm *algorithm = controller->GetAlgorithm("awb");
AwbAlgorithm *awb_algorithm = dynamic_cast<AwbAlgorithm *>(algorithm);
if (awb_algorithm) {
    awb_algorithm->SetMode("sunny");
}
```

5.8.10. Parameters

Name	Default	Description
`bayes`	1	Whether to use the Bayesian algorithm, or the simplified Grey World algorithm. Note that if the CT Curve is not defined the implementation will default back to Grey World in any case.
`frame_period`	10	Only run the AWB calculation every time this many frames have elapsed.
`speed`	0.05	IIR filter speed determining how quickly the image adapts towards the most recent results of the AWB calculation. A speed of zero effectively disables AWB, and a value of 1 causes the AWB calculation results to be followed immediately with no damping.
`startup_frames`	10	At startup, run the AWB algorithm as often as possible (ignoring `frame_period`) for this many frames. During these frames the `speed` parameter is also treated as having the value 1. This speeds up initial convergence when a camera application is opened.
`convergence_frames`	3	When queried by the *GetConvergenceFrames* method, the algorithm will return this value as the number of frames to drop.
`ct_curve`	*Optional*	Defines the functions $r = r(T)$ and $b = b(T)$. The numbers in this list always appear in threes. The first value of the triple is the colour temperature (T), the second is the r value and the third is the b value. The triples should be listed in increasing order of colour temperature.
`priors`	*Optional*	A list of illuminant priors for different lux levels. For each lux level the prior is a list of numbers appearing in pairs, where the first is the colour temperature and the second is the value of the log likelihood. If omitted, the priors are simply assumed to be identically zero everywhere.
`modes`	*Required*	A list of AWB modes giving a name (a character string) and the “lo” and “hi” range in terms of colour temperature on the CT curve that must be searched.
`min_pixels`	16	Minimum number of pixels in a statistics region for that region to be used in the AWB calculations.
`min_G`	32	Minimum average green value in a statistics region (out of a 16-bit dynamic range) for the region to be used in the AWB calculations.
`min_regions`	10	Minimum number of usable statistics regions for the AWB algorithm to run (otherwise the target red and blue gains will be left unchanged).
`delta_limit`	0.15	Limit to the value of the squared colour error term in the log likelihood expression; stops very non-neutral colours being excessively penalised.
`coarse_step`	0.2	Governs the size of the steps taken in the coarse search phase of the algorithm.
`transverse_pos`	0.01	How far off the CT curve the fine search is allowed to go, in the direction of less green illuminants.
`transverse_neg`	0.01	How far off the CT curve the fine search is allowed to go, in the direction of more green illuminants.
`sensitivity_r`	1.0	Calibration of the red sensitivity of the sensor being used compared to the one that was used for AWB tuning. This value cannot be set by the tuning process and would need some form of measurement for every sensor individually. Normally it lies close enough to 1.0 to make little difference.
`sensitivity_b`	1.0	Calibration of the blue sensitivity of the sensor being used compared to the one that was used for AWB tuning. This value cannot be set by the tuning process and would need some form of measurement for every sensor individually. Normally it lies close enough to 1.0 to make little difference.
`whitepoint_r`	0.0	The whitepoint for the Δr coordinate.
`whitepoint_b`	0.0	The whitepoint for the Δb coordinate.

5.8.11. Metadata Dependencies

The AWB control algorithm requires the presence of the `LuxStatus` metadata from the Lux control algorithm.

The AWB control algorithm writes an `AwbStatus` object (`controller/awb_status.h`) into the image metadata.

5.8.12. Extensions

There are many relatively minor improvements and tweaks we could make to the algorithm described here. However, we have predominantly chosen to keep the approach simple and avoided anything that would complicate either the algorithm itself or the tuning process it would require, even at the cost of some modestly improved performance. We discuss just a couple of such ideas here.

5.8.13. Uniform Colour

One possible weakness of the algorithm is that it could be influenced excessively by large objects of a uniform but non-grey colour. In Bayesian terms, if you have already seen something bright red in an image, then it's more likely for neighbouring parts of the image to have the same colour as it might just be the same object. This obviously strikes at that assumption of independence that we made earlier, and such ideas are not difficult to incorporate. When we calculate the sum of squares term in the log likelihood expression for a particular patch in the image (*viz.* $\Delta r^2 + \Delta b^2$) we could also calculate a similar term measuring the distance from (some of) the neighbours' colours, and then simply pick the smallest.

Of course, experimentation would be required with a reasonable corpus of test images to determine the precise details, not least the relative strength of this effect overall and how it might vary.

5.8.14. Non-grey Colours

Another fairly clear improvement might be to "look for" colours other than grey in an image. For instance, it might prove beneficial to look for things like blue sky, green grass and even skin tones. At one level this is not difficult either. We compute the error term as before, but also compute the distance from these other measured colours, and again take the smallest. However, this does end up complicating the algorithm and can make the tuning process more difficult. In particular:

- These additional target colours would need to be measured, probably in the colourspace of the camera, and this causes additional calibration/tuning problems. Or you could try to measure them in sRGB which might save remeasuring them all the time - though then the AWB algorithm will implicitly include a colourspace conversion, which is all starting to get a bit circular.
- These extra colours are not likely to be present in all situations. For example blue sky is unlikely to feature in low light images, and matching it erroneously could throw the whole algorithm off. So there would have to be additional *prior knowledge* for each target colour expressing its effect in the known conditions.
- Finally, we have ended up with quite a lot of extra effects, parameters and new kinds of *priors* to calibrate. Determining appropriate values is non-trivial, and depending on the details, may have to be repeated for every new camera module.

5.9. AGC/AEC (Automatic Gain Control / Automatic Exposure Control)

5.9.1. Name

`rpi.agc`

5.9.2. Overview

The AGC/AEC algorithm controls the shutter time of the image sensor, the analogue gain applied by the image sensor and the digital gain applied by the ISP. Normally we refer simply to the "AGC algorithm" but it encompasses all these aspects. The metering process uses both a matrix of rectangular regions spread across the image and constraints derived from an intensity histogram of pixels in the image, as explained in the following sections.

Multi-channel AGC

Normally AGC/AEC operates equally on every frame from the camera, which is the normal "single channel" mode of operation. Optionally, AGC can be instructed to operate independently, with different parameter settings, on different frames. For example,

you might want to meter odd and even numbered frames differently, perhaps as part of a more sophisticated multi-exposure camera application.

5.9.3. Region-based Metering

The image is divided up into a number of distinct regions and an average luminance value for the whole image is calculated as

$$Y = \frac{\sum_{i \in \text{ regions}} w_i Y_i}{\sum_{i \in \text{ regions}} w_i}$$

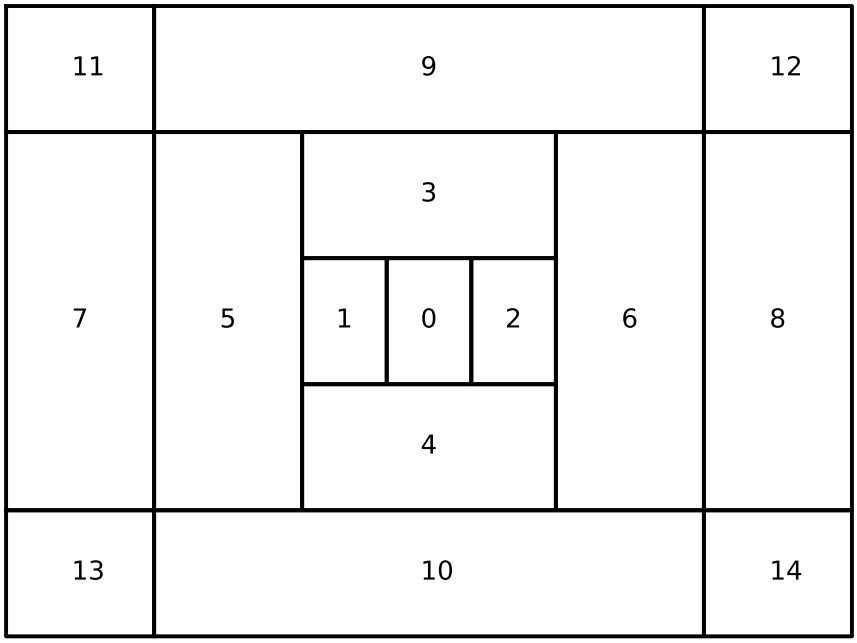
where Y_i is the luminance of region i and w_i is the weight associated with that region and which is defined in the JSON camera tuning file.

This value is driven towards a target value `y_target` , also given in the tuning file. The weights can be adjusted to create various forms of average, centre-weighted or spot metering.

Pi 4 and earlier devices

On Pi 4 and earlier devices, the image is divided up into regions as shown in .

Figure 9.
AGC metering regions.



Pi 5

On Pi 5, the image is divided up into 16x16 equally sized regions covering the whole image. The list of weights in the tuning file is therefore much larger, being a 16x16 matrix of numbers. Each weight must be an integer from 0 (meaning that the region is ignored) to 15 (for the most important regions).

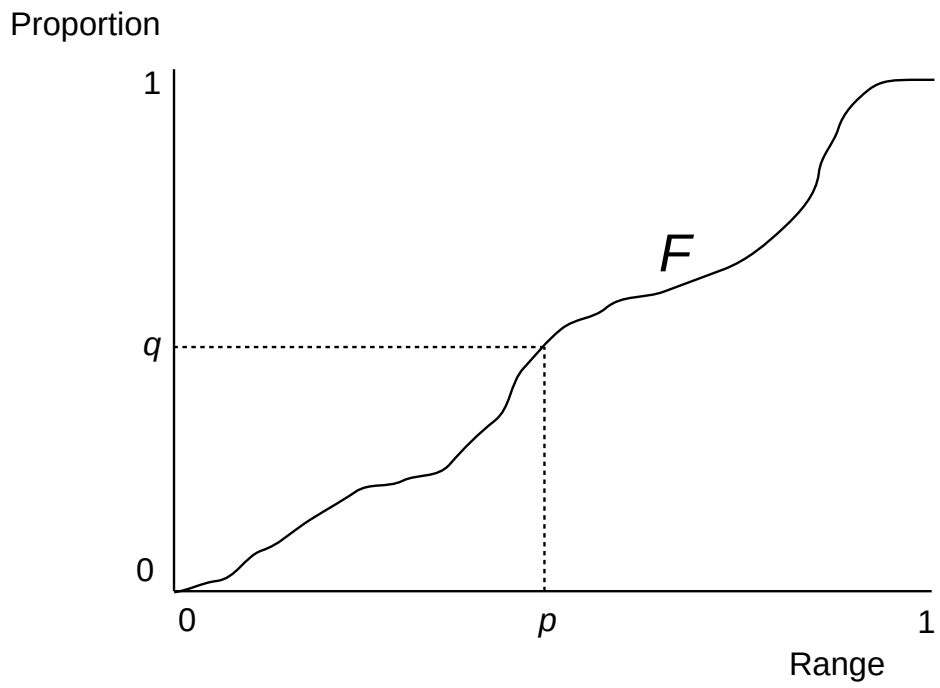
5.9.4. Histogram Constraints

Beyond controlling the weighted luminance value, the algorithm also uses a histogram, collected by the ISP for this frame (Pi 5) or on an earlier frame (Pi 4 and earlier devices), to impose various other constraints on the final target exposure value.

We assume basic familiarity the the idea of an intensity histogram for an image. We define $F(p)$ to be the cumulative frequency of the histogram, and further normalise it so that the input value p is 0 at the bottom of the histogram and 1 at the very top. We also normalise the result so that $F(p) = 0$ indicates there are no pixels below p and $F(p) = 1$ means than all pixels lie below p. Further,

we adopt the convention that pixels are spread evenly throughout the "bin" in which they lie, so that $F(p)$ is a continuous (in fact, monotonically increasing) function.

Figure 10.
Cumulative frequency: a proportion q of the pixels lie below the value p



Observe that we always have $F(0) = 0$ and $F(1) = 1$.

Next we recap the notion of a *quantile*, which gives us the point $p = Q(q)$ in the range such that a proportion q of the pixels lie below p. Some quantiles are familiar to us, for example $Q(0.25)$, $Q(0.5)$ and $Q(0.75)$ are known as, respectively, the lower quartile, the median and the upper quartile of a distribution. We note further that F and Q are in some sense inverses, as $Q(F(p)) = p$, though this breaks down when F has flat regions (that is, there are no pixels in a part of the range) and Q becomes non-single-valued.

Q being ill-defined explains why quantiles are not ideal for metering as they are. Imagine, for instance, a picture containing only black and white pixels (an image of a newspaper page, perhaps). The median pixel will be either white or black, and will flick suddenly between them when only a single extra pixel pushes the number of white pixels either above or below the 50% mark. Any metering scheme based on this will be subject to occasional violent oscillations. Instead the Raspberry Pi algorithm uses the concept of the *inter-quantile mean* instead. Here we define $I(q_{lo}, q_{hi})$ to be the mean of all the pixels between the q_{lo} and q_{hi} quantiles. Observe how the possibly ill-defined nature of these locations in the pixel range become irrelevant as there are then no pixels to average there. We refer to $q_{hi} - q_{lo}$ as the *width* of the inter-quantile mean. We note that

- A narrow width gives precise control of a particular place in the histogram.
- A large width gives a stable response in the location of the inter-quantile mean.
- In practice we of course have to find some kind of happy medium.

Finally, our histogram constraints comprise therefore an inter-quantile mean (specifying values for q_{lo} and q_{hi}), a target Y value for it, and an indication of whether this is to be a *lower* or an *upper* bound for our final target exposure value. We illustrate this with a pair of examples.

1. $I(0.98, 1) = 0.5$, lower bound. Here we are saying that the top 2% of the histogram must lie at or above 0.5 in the pixel range (metering all happens before gamma, so this is a moderately bright value). Or in short, we are requiring "some of the pixels to be reasonably bright". This is a good strategy for snowy or beach scenes, also for images of documents. It actually raises the exposure of the whole scene; without it, the entire image would look dull grey which is undesirable in such circumstances. On the other hand when, as is usually the case, there is an abundance of bright and dark pixels, this constraint has no effect, so it is often reasonable to apply this constraint all the time.
2. $I(0.98, 1) = 0.8$, upper bound. This requires the top 2% of pixels to lie at or below 0.8 in the pixel range (actually a very bright value post-gamma). Effectively this lowers the exposure in order to stop pixels saturating, and is effective in metering for highlights or perhaps as part of an HDR metering strategy.

A number of these constraints can be grouped together as a named *constraint mode*, in which they are applied in sequence one after another. Applications are able to choose which of the available constraint modes in the JSON file they wish to use.

Differences between Pi 4 and Pi 5

On a Pi 4 (or earlier) the histogram counts every pixel just once. On a Pi 5, however, each pixel is counted w_i times, where w_i is the metering weight for the region containing the pixel. The Pi 4 also uses separate R, G and B histograms meaning that histogram constraints really apply to the green pixels, whereas the Pi 5 has access to a proper Y histogram.

5.9.5. Channel Constraints

There is one further type of constraint that gets applied in the case of multi-channel AGC. Here, we can constrain one channel to be, for example, “more than 2 times the exposure of another channel” or “less than 4 times the exposure of another channel”. For example

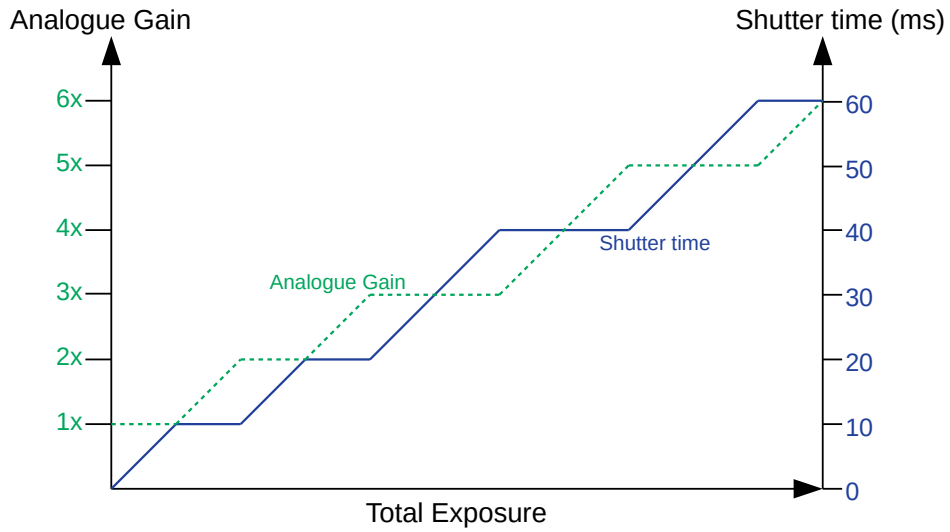
```
Code
...
"channel_constraints": [
    {
        "bound": "UPPER",
        "channel": 2,
        "factor": 8
    }
]
...
```

would constrain a channel to have no more than 8 times the exposure of channel 2. Use `LOWER` to enforce a lower bound. Several such constraints may be included in the `channel_constraints` list, and they are applied in the order listed.

5.9.6. Exposure Modes

The region-based averages and histogram constraints determine the total exposure of the final scene, however they do not determine how to divide this up between the shutter time and the analogue gain. This is the job of the *exposure mode*. Each named exposure mode (in the JSON file) consists of a list of shutter times and analogue gains. First the desired analogue gain is set to 1 and the shutter time is allowed to ramp to the first value in the list. Thereafter, the analogue gain is allowed to ramp to the first value in its list. The procedure then simply repeats with the second value in each of the lists.

Figure 11.
An example exposure profile - analogue gain and shutter time increase one after the other.



The maximum possible total exposure is determined by the final values in these two lists. Note how

- The maximum shutter time can be limited to avoid dropping the frame rate (assuming the camera driver supports this behaviour).

- If the camera driver does not support an analogue gain as high as listed in the exposure modes, the AGC algorithm will simply make up the difference with digital gain.
- For any total exposure value (less than the maximum), there is only a single pair of analogue gain and shutter time values that achieves it.

5.9.7. Implementation

The implementation of the AGC algorithm can be found in `controller/rpi/agc_channel.cpp`. We note the following.

- Y targets are generally defined by piecewise linear functions, making it possible to vary the Y target with estimated lux level.
- Mostly the algorithm can be regarded as simply adjusting the total exposure of the images, filtering the requested values over time to prevent sudden changes. When it needs to reduce the exposure substantially, it may elect to cut the camera exposure more rapidly, whilst (temporarily) hiding any sudden changes from observers using increased digital gain. (Please refer to the source code for more details.)
- The algorithm is able to avoid exposure times that will result in flicker under artificial lighting. It would have to be given the period (in microseconds) of the lighting cycle.
- The implementation allows for the application to fix the shutter time and/or the analogue gain (for example, for fixed ISO exposures).
- The Raspberry Pi AGC algorithm does not currently handle flashes as we have no camera modules or boards that incorporate them.
- The Raspberry Pi AGC algorithm does not currently handle variable apertures as we have no camera modules that feature them.
- Our `Histogram` class, which provides histogram-related methods, is implemented in `controller/histogram.cpp`, though it does not in fact normalise the input and output ranges of the histogram as we did in the theoretical discussion above.

5.9.8. External API

The `Agc` class is derived from the `AgcAlgorithm` class. As such it defines the following publicly accessible methods. These APIs apply to all AGC channels, unless there is an explicit `channel` parameter.

AgcAlgorithm class API
`GetConvergenceFrames()` This returns the number of frames the AEC/AGC algorithm recommends be dropped, while the AGC converges, when the camera and the AGC algorithm are started from scratch.
`SetEv(unsigned int channel, double ev)` Sets the exposure compensation according to a linear scale. So for an extra stop, use `ev` = 2. For an extra half stop, use `ev` = 1.414. The `channel` identifies the AGC channel the setting applies to, which will be zero if running in single channel mode.
`SetFlickerPeriod(double flickerPeriod)` Set the lighting cycle period in microseconds. Set to zero to disable flicker avoidance.
`SetFixedShutter(unsigned int channel, double fixedShutter)` Sets a fixed shutter time, meaning only analogue (and digital) gain can be used to affect exposure. Set to 0 to cancel fixed shutter operation. The `channel` identifies the AGC channel the setting applies to, which will be zero if running in single channel mode.
`SetMaxShutter(double maxShutter)` Sets maximum allowable shutter speed to use. This is typically based on the application requested framerate. Set to 0 to disable any limitation on the maximum shutter speed.
`SetFixedAnalogueGain(unsigned int channel, double fixedAnalogueGain)` Sets a fixed analogue gain value, meaning only shutter time can be varied to affect exposure. Set to 0 to cancel fixed analogue gain operation. The `channel` identifies the AGC channel the setting applies to, which will be zero if running in single channel mode.
`SetMeteringMode(std::string const &meteringModeName)` Choose which metering mode to use from the JSON file.
`SetExposureMode(std::string const &exposureModeName)` Choose which exposure mode to use from the JSON file.
`SetConstraintMode(std::string const &constraintModeName)` Choose which (histogram) constraint mode to use from the JSON file.
`SetActiveChannels(const std::vector<unsigned int> &activeChannels)` Separate instances of AGC are run repeatedly in the exact order given in this list. For example, passing the values `{1, 2}` will alternate between channels 1 and 2 (which would have to be defined in the tuning file).

An example use of AGC methods:

```
Code
Algorithm *algorithm = controller->GetAlgorithm("agc");
AgcAlgorithm *agc_algorithm = dynamic_cast<AgcAlgorithm *>(algorithm);
if (agc_algorithm) {
    agc_algorithm->SetEv(1.414);
    agc_algorithm->SetExposureMode("sport");
}
```

5.9.9. Parameters

Name	Default	Description
`startup_frames`	10	For this many frames the AGC algorithm runs with maximum speed (1.0) so as to adapt quickly to the initial conditions.
`convergence_frames`	6	When queried by the *GetConvergenceFrames* method, the algorithm will return this value as the number of frames to drop, unless both shutter and gain values have been set (when there is no convergence required).
`desaturate`	*1*	When set, AGC will sometimes decide to lower the camera exposure to meter more effectively for saturated image regions, whilst hiding this sudden reduction from the application by applying digital gain to compensate. This behaviour can be suppressed by setting the value zero.
`metering_modes`	*Required*	A list of metering modes. Each mode specifies a name and 15 numbers, which are the weights for the region-based AGC calculation.
`exposure_modes`	*Required*	A list of exposure modes. Each mode specifies a name and a list of of *shutter* and *gain* values which define how the total exposure time is to be broken up between shutter time and analogue gain.
`constraint_modes`	*Required*	A list of constraint modes. Each mode specifies a name and a list of histogram constraints, involving the bound type (UPPER or LOWER), q values for an inter-quantile mean, and a Y target value for it.
`y_target`	*Required*	Piecewise linear function defining a Y target value (out of a maximum value of 1.0) for the region-based metering.
`base_ev`	1.0	Optional parameter that linearly scales the final target exposure value of the algorithm. Useful for tuning if ever it is desired for the AGC to be globally a bit brighter (> 1.0) or darker(< 1.0) purely for reasons of taste.
`speed`	0.2	Speed of adaptation of the AGC algorithm. A larger value (with a limit of 1.0) makes the algorithm adjust the camera's exposure more quickly to the scene conditions.

Single and Multi-channel operation

The parameters above can be listed directly in the `rpi.agc` section of the camera tuning file for single channel operation, as shown below.

```
Code
...
{
    "rpi.agc":
    {
        # parameters here (which define channel 0)
    }
},
...
```

For multi-channel operation, the `channels` keyword should be used, followed by a list of parameter groups, each of which defines a single channel. There is no specific limit to the number of available channels. Although entirely optional, there is a convention of putting a `comment` into each channel that briefly describes it.

```
Code
...
{
    "rpi.agc":
    {
        "channels": [
            {
                "comment": "This is normal AGC metering",
                # parameters for channel 0
            },
            {
                "comment": "This is a modified AGC configuration",
                # parameters for channel 1
            }
        ]
},
...
```

Defining multi-channel AGC (with the `channels` keyword) but listing parameters for only a single channel (channel 0) is identical to the single-channel example that we saw previously (without the `channels` keyword). By default, when AGC starts, it runs using only channel 0, even when there are more available.

5.9.10. Metadata Dependencies

The AGC control algorithm requires the presence of

1. The `DeviceStatus` metadata (which should always be available from the camera).
2. The `LuxStatus` metadata from the Lux control algorithm.
3. The `AwbStatus` metadata from the AWB algorithm. Note that it is therefore best to list the AGC algorithm in the JSON file after the AWB algorithm (there is no reverse dependence of the AWB on the AGC algorithm).

The AGC control algorithm writes an `AgcStatus` object (`controller/agc_status.h`) into the image metadata in its `process` method, and an `AgcPrepareStatus` object (same file) in its `prepare` method.

5.10. ALSC (Automatic Lens Shading Correction)

5.10.1. Name

`rpi.alsc`

5.10.2. Overview

Lens shading, also known as *vignetting*, is the problem where light does not fall equally on all parts of the sensor, particularly near the edges. It causes obvious darkening but also discolouration to parts of the image, as the effect is not identical in all three colour channels.

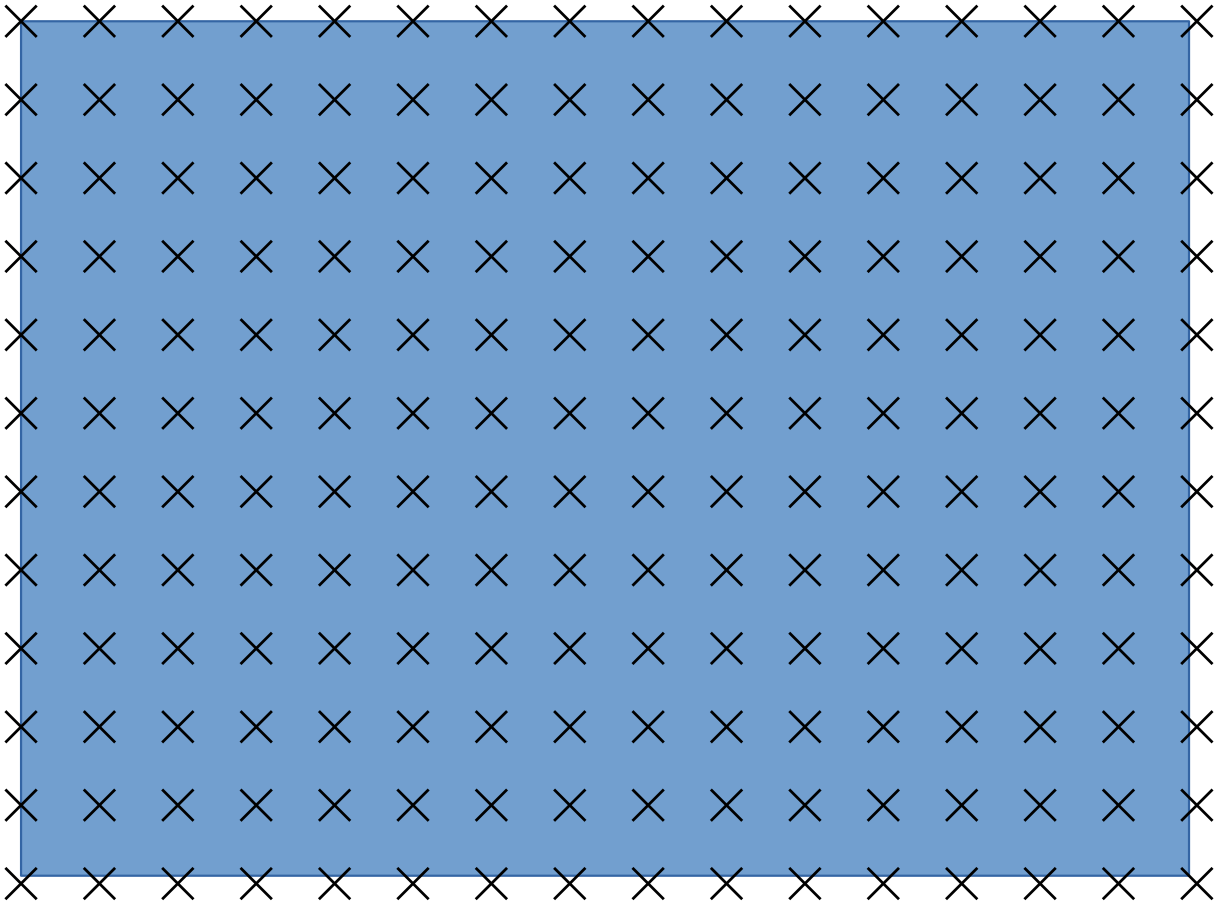
Figure 12.
Lens shading: no correction (left), luminance correction only (middle), full correction (right).



We see how the uncorrected image (left) shows obvious darkening towards the edges. The middle image corrects this luminance fall-off, but we still see the colours shifting towards the image edge. Whilst the very centre isn't too bad we clearly see progressively more of a cyan cast as we move away. The final image on the right shows how the image should look - there is no (or minimal) luminance fall-off, and the colour shift has been corrected.

Lens shading is corrected by creating a spatial 2-dimensional table of gain values that covers the image in question. The gains are located at regular horizontal and vertical intervals and pixels are corrected by interpolating a local gain value from the four nearest values on the grid. We create a separate table for each of red, green and blue so as to be able to correct the varying colours caused by the lens shading effect.

Figure 13.
A grid of 16x12 lens shading gains. The grid may extend slightly beyond the image to ensure every pixel is covered. On a Pi 5, the grid consists of 32x32 regions.



It is generally found that while luminance correction is beneficial it is usually not necessary to correct it fully. Firstly we are quite accustomed to seeing images with at least a little darkening in the periphery and secondly, full correction can boost noise more than we might like. However, colour shading problems are generally viewed as more serious as they cause very obvious colour shifts and casts (especially when the background of an image is very uniform), and there is no real upside to not fixing them.

The Raspberry Pi ALSC algorithm separates the luminance and colour shading problems. We always correct any colour shading artifacts fully, and then a proportion of the luminance fall-off according to preference. We calculate separate tables for the red and blue channels to correct the colour shading (and set the green table to unity everywhere) and then post-multiply each of these by a single luminance correction table.

ALSC tables and parameters would normally be calculated automatically from a set of calibration images by the Camera Tuning Tool.

5.10.3. Luminance Shading

Luminance shading is measured straightforwardly from a set of calibration images. Whilst the algorithm calculates only a single average luminance correction table, there would be no difficulty in extending this to a more complex scheme should we wish.

5.10.4. Colour Shading

This is by far the more difficult part of the problem. We take a hybrid approach to the problem, calculating the red and blue colour correction tables partly by calibration and partly using an adaptive procedure.

Historically the preferred method for colour shading correction has been to measure separate correction tables for each of red and blue under different colour temperature illuminants. Then the correct tables are chosen (or interpolated) according to the current colour temperature being reported by the AWB algorithm. There are a couple of problems with this.

1. The nature of the colour shading actually varies with the whole spectrum of the illuminant, not just its apparent colour temperature. Using only the colour temperature based tables will invariably leave some illuminants with residual colour shading problems (this can be a particular problem when dealing with mixtures of incandescent, fluorescent or other types of light).
2. There are also variations between different camera modules. Whilst we may tune the colour shading correction tables with a particular reference (or "golden") sensor, other sensors in use out in the field may respond slightly differently, also causing residual colour shading problems (unless we go to great expense and calibrate every single module).

Our solution is to fix "as much as we can" using the colour temperature based calibration method, and then further to apply the adaptive procedure to "mop up" some of the residual problems. The adaptive procedure is calibrated so as not to be more aggressive than is necessary.

5.10.5. Colour Temperature based calibration

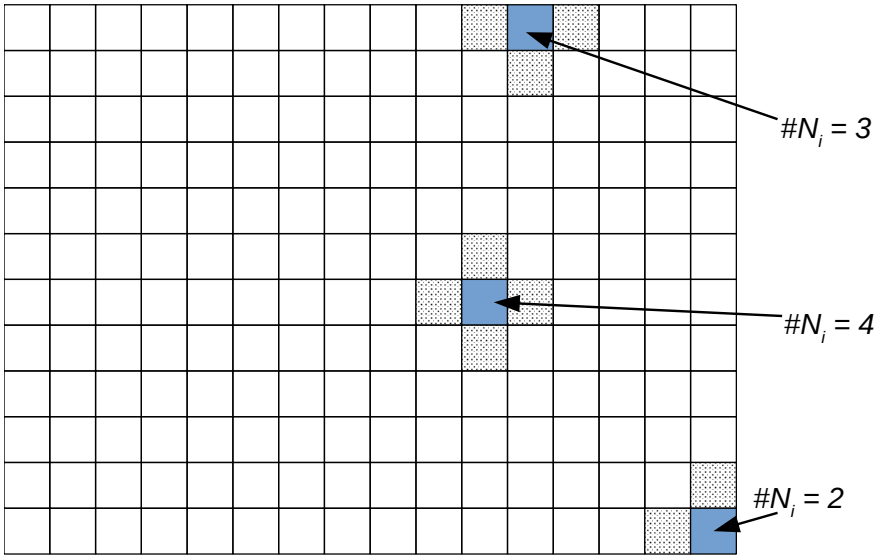
Here we simply list tables of G/R (or G/B for the blue tables) for our calibration images. We may list a separate table for each colour temperature that we have. As the algorithm runs it will interpolate as necessary so as to create an appropriate estimate of a table for the current illuminant. These tables are then applied to the *statistics* information from the ISP, so that the adaptive algorithm, which runs next, only makes the additional changes required on top of the calibration-based fixes.

5.10.6. Adaptive ALSC Algorithm

Without loss of generality we shall assume we are working with the red colour channel. Exactly the same process can be followed subsequently for blue (*mutatis mutandis*). We also denote each of the grid regions (or cells) with a number $i, 0 \leq i < N$, where N is the total number of grid cells ($16 \times 12 = 192$ in the case of the VC4 platform, or $32 \times 32 = 1024$ on a Pi 5).

We start by defining a chrominance statistic for each grid cell so that $C_i = R_i/(G_i + k)$, using a small constant k to avoid numerical instabilities and potential division by zero, and also a neighbourhood $\mathcal{N}_i$ consisting of the north, south, east and west neighbours of grid cell i. Writing $\#\mathcal{N}_i$ to denote the number of neighbours of a grid cell we note that $\#\mathcal{N}_i = 4$ for all interior cells, 3 for edge cells and 2 for the four corner cells.

Figure 14.
Cells and their neighbours in a 16x12 grid. The same kinds of definitions hold for the 32x32 grid on a Pi 5.



Now, the problem at hand is to find multipliers λ_i so that $\lambda_i C_i$ is the same "in places where it looks like it should be". To cope with this somewhat nebulous phrase we introduce a matrix of weights w_{ij} defined for $j \in \mathcal{N}_i$ and zero otherwise. Notice that w_{ij} forms a large (192 x 192, or 1024 x 1024 on a Pi 5) but sparse matrix with only $\#\mathcal{N}_i$ non-zero entries on each row. Where non-zero, w_{ij} will be calculated by comparing C_i and C_j. When they are very similar then w_{ij} takes a value near one; when they are very different, the value goes to zero. It could be implemented by a Gaussian-style $\left(e^{-x^2}\right)$ function.

Now we can define a constraint saying that $\lambda_i C_i$ should be the same as its neighbours, taking account of the defined weights. Thus:

$$\lambda_i C_i = \frac{\sum_{j \in \mathcal{N}_i} w_{ij} \lambda_j C_j}{\sum_{j \in \mathcal{N}_i} w_{ij}}$$

There is obviously a danger here that all the w_{ij} can go to zero, so in practice we add another term weighted by the small value ε saying that "when the neighbours don't help, simply use the average of the neighbour's λ values". Hence:

$$\lambda_i C_i = \frac{\sum_{j \in \mathcal{N}_i} w_{ij} \lambda_j C_j + \frac{\varepsilon}{\#\mathcal{N}_i}\left(\sum_{j \in \mathcal{N}_i} \lambda_j\right) C_i}{\sum_{j \in \mathcal{N}_i} w_{ij} + \varepsilon}$$

Rearranging a little, we end up with a linear system of N equations:

$$\left(\varepsilon + \sum_{j \in \mathcal{N}_i} w_{ij}\right) C_i \lambda_i - \sum_{j \in \mathcal{N}_i} \left(w_{ij} C_j + \frac{\varepsilon}{\#\mathcal{N}_i} C_i\right) \lambda_j = 0$$

Rewriting these coefficients as a matrix M and the λ_i as a vector λ we finally have

$$M\lambda = 0$$

5.10.7. Solution

Astute readers will observe that, as there is no particular reason why the matrix M is singular, then there is only a single solution to this system, namely $\lambda = 0$. In practice, however, the equations lend themselves to the *Gauss-Seidel* method wherein λ_i is updated from the i^{th} equation according to:

$$\lambda_i = \frac{\sum_{j \in \mathcal{N}_i} \left(w_{ij} C_j + \frac{\varepsilon}{\#\mathcal{N}_i} C_i\right) \lambda_j}{\left(\varepsilon + \sum_{j \in \mathcal{N}_i} w_{ij}\right) C_i}$$

Note how, informally, $\lambda_i C_i$ is replaced by a convex combination of its neighbours; therefore the iterations do not run off towards zero but remain bounded, iterating to a place of lower overall colour error.

5.10.8. Implementation

The ALSC algorithm is implemented in `controller/rpi/alsc.cpp`.

- The statistics are in fact obtained with the colour shading correction already applied. Therefore the algorithm has to start by dividing out its estimate of the previous correction before continuing (the function `copy_stats`).
- The calibrated tables are expected to be for the full field of view sensor frame. When operating in a cropped mode these tables need to be cropped and resampled (see `resample_cal_table`).
- The Gauss-Seidel method is applied with a small amount of over-relaxation so as to reach a point of low colour error more quickly (`gauss_seidel2_SOR`).
- We apply a normalisation to the λ values, such that $\min_i \lambda_i = 1$. This is because it is only their relative values that matter and we do not wish to add any unnecessary extra gain into the image (or indeed apply any gains less than unity).
- The main part of the algorithm - constructing the matrix M and the Gauss-Seidel iterations - runs asynchronously, and only every few frames. Other than that the implementation closely follows the method and nomenclature in this document.
- Lens shading algorithms do not have any additional public methods (at this time).

5.10.9. Parameters

Name	Default	Description
`frame_period`	12	The ALSC algorithm runs once per this many input frames. Note that new tables are generated on every frame, being filtered gradually towards the results of the most recent run of the ALSC computation.
`startup_frames`	10	For this many frames the ALSC algorithm runs as often as possible (ignoring `frame_period`) so as to adapt quickly to the initial conditions.
`speed`	0.05	Coefficient of the IIR filter that adapts the tables, on every frame, towards the latest ALSC result. A value of 1 would adopt the most recently calculated results immediately.
`sigma`	0.1	Default value for the standard deviation of colour differences (R/G or B/G), used in calculating the weights w.
`sigma_Cr`	`sigma`	Standard deviation of red channel colour differences (R/G).
`sigma_Cb`	`sigma`	Standard deviation of blue channel colour differences (R/G).
`min_count`	10	Number if pixels in a statistics region for that region's data to be considered meaningful.
`min_G`	800	Average green value (out of a 16-bit dynamic range) in a statistics region for that region's data to be considered meaningful.
`n_iter`	28	Maximum number of Gauss-Seidel iterations for perform.
`omega`	1.3	Amount of over-relaxation to apply in Gauss-Seidel iterations. A value greater than 1 implies (increasing) over-relaxation; values less than 1 would imply under-relaxation.
`calibrations_Cr`	*Optional*	Tables, one for each colour temperature listed, giving the gains required to correct the red colour channel for that colour temperature in a 16x12 spatial grid. If omitted, the algorithm will behave as if these tables contained unity values everywhere.
`calibrations_Cb`	*Optional*	Tables, one for each colour temperature listed, giving the gains required to correct the blue colour channel for that colour temperature in a 16x12 spatial grid. If omitted, the algorithm will behave as if these tables contained unity values everywhere.
`luminance_strength`	1.0	Proportion of the full luminance correction to apply, so that zero means no luminance correction and 1 means 100% correction.
`luminance_lut`	*Optional*	16x12 spatial luminance correction table. This table gives gains that will be applied to all colour channels for full, or 100% luminance correction.
`default_ct`	4500	Colour temperature (in kelvin) to be assumed if no AWB metadata is found.
`threshold`	1e-3	Gauss-Seidel iterations stop once the step sizes in the solution are all less than this value.

5.10.10. Metadata Dependencies

The AGC control algorithm requires the presence of the `AwbStatus` metadata from the AWB algorithm. Note that it is therefore best to list the ALSC algorithm in the JSON file after the AWB algorithm (there is no reverse dependence of the AWB on the AGC algorithm).

The ALSC control algorithm writes an `AlscStatus` object (`controller/alsc_status.h`) into the image metadata.

5.11. Contrast

5.11.1. Name

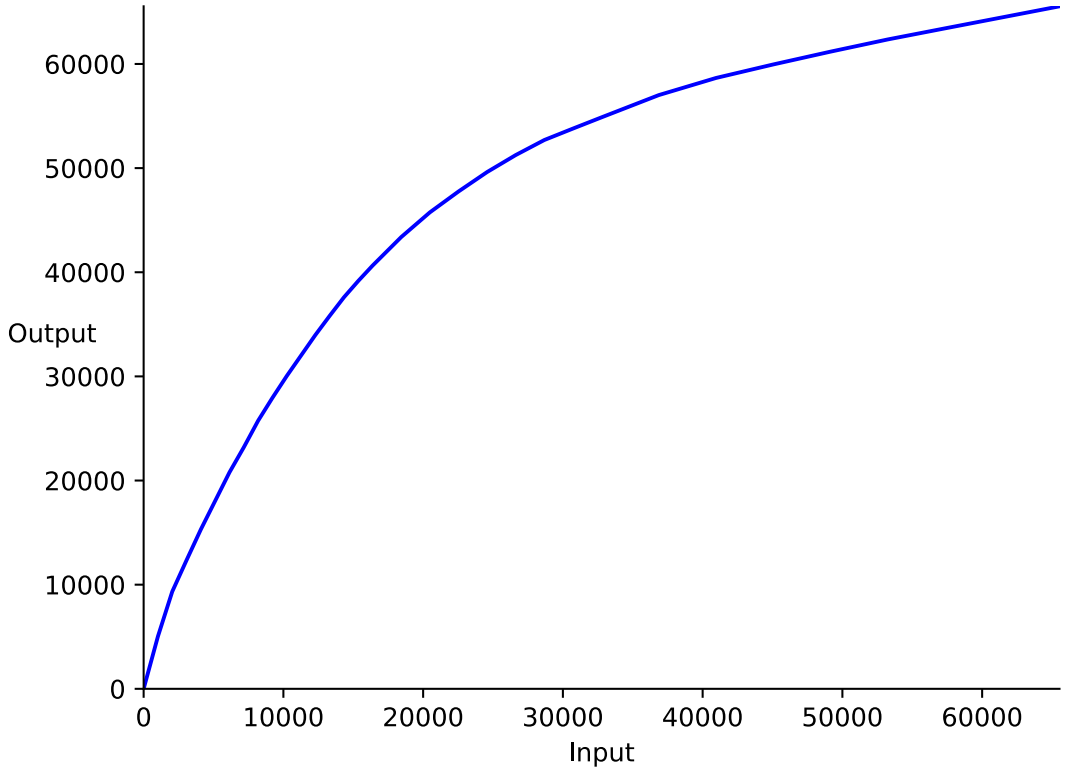
`rpi.contrast`

5.11.2. Overview

The Contrast Control Algorithm is responsible for calculating a gamma curve with which to program the ISP hardware. In the first instance, the JSON tuning file may supply a gamma curve, and applications may allow users manually to change contrast and

brightness (as is customary with monitors and TVs). Additionally, the algorithm has some ability to watch the image histogram being reported by the ISP and alter the gamma curve slightly so as to enhance the global contrast.

Figure 15.
An example gamma curve with 16-bit inputs and outputs.



Note how the gamma curve is more properly an *inverse* gamma curve, looking more like $y = x^{1/2.2}$ than $y = x^{2.2}$. Furthermore, the "theoretical 2.2 gamma" is rarely used in practice as it gives rise to rather washed-out low contrast images.

5.11.3. Adaptive Contrast Enhancement

Adaptive Contrast Enhancement uses the image histograms collected by the ISP. It works by selecting a particular *quantile* near the bottom of the histogram and moving it to a particular place in the output pixel range. A similar calculation is performed for a point near the top of the histogram. Normally we use this to ensure the resulting images have sufficiently strong contrast.

The algorithm also allows brightness and contrast to be adjusted globally. After applying the gamma curve to the pixel data, the following transformation is applied.

$$pixel_{out} = (pixel_{in} - 32768)*contrast + 32768 + brightness$$

(note: pixel values are always treated as being unsigned 16-bit numbers).

5.11.4. External API

The `Contrast` class is derived from `ContrastAlgorithm` . This means it implements the following additional public methods.

ContrastAlgorithm class API
`setBrightness(double brightness)` Sets a fixed brightness offset (a pixel value out of a 16-bit range).
`setContrast(double contrast)` Sets a fixed contrast value. A value less than one decreases the contrast, a value greater than one increases the contrast.

Code

```
Algorithm *algorithm = controller->GetAlgorithm("contrast");
```

```
ContrastAlgorithm *contrast_algorithm = dynamic_cast<ContrastAlgorithm *>(algorithm);
if (contrast_algorithm) {
    contrast_algorithm->setBrightness(2000);
    contrast_algorithm->setContrast(1.1);
}
```

5.11.5. Parameters

Name	Default	Description
`startup_frames`	6	The first few frames emerging from the ISP can be rendered entirely as black, in case algorithms like AEC/AGC are still converging (and application code is not hiding them).
`ce_enable`	1	Whether to enable adaptive contrast enhancement. If set to zero, the gamma curve is used exactly as specified (subject to user brightness/contrast settings).
`lo_histogram`	0.01	Quantile near the bottom of the histogram to be moved when `ce_enable` is set.
`lo_level`	0.015	Level (out of a maximum of 1.0) in the range where the current `lo_histogram` quantile will be moved.
`lo_max`	500	Maximum amount (out of 16 bits) by which the low quantile may be moved.
`hi_histogram`	0.95	Quantile near the top of the histogram to be moved when `ce_enable` is set.
`hi_level`	0.95	Level (out of a maximum of 1.0) in the range where the current `hi_histogram` quantile will be moved.
`hi_max`	2000	Maximum amount (out of 16 bits) by which the high quantile may be moved.
`gamma_curve`	*Required*	Gamma curve. A list of alternating input and output points, each given out of a 16-bit range. There must be an even number of points in this list.

5.11.6. Metadata Dependencies

The Contrast Control Algorithm does not require any other metadata.

The Contrast Control Algorithm writes a `ContrastStatus` object (`controller/contrast_status.h`) into the image metadata.

5.12. CCM (Colour Correction Matrices)

5.12.1. Name

`rpi.ccm`

5.12.2. Overview

The CCM (Colour Correction Matrix) Algorithm selects an appropriate CCM for the prevailing conditions to convert the camera's notion of RGB into our standard version of RGB according to

$$\begin{pmatrix} R_{standard} \\ G_{standard} \\ B_{standard} \end{pmatrix} = C \begin{pmatrix} R_{camera} \\ G_{camera} \\ B_{camera} \end{pmatrix}$$

where C is a 3x3 matrix. Note that the rows of C will normally sum to 1 (otherwise grey pixels will not be preserved).

Calibration is usually performed using the Camera Tuning Tool, and consists of a list of colour temperatures, each followed by a 3x3 matrix written out as 9 values in normal reading order. During operation, the algorithm will estimate a CCM - interpolating as necessary - using the current estimated colour temperature as reported by the AWB algorithm.

In the following example colour matrix

$$\begin{pmatrix} 1.80439 & -0.73699 & -0.06739 \\ -0.36073 & 1.83327 & -0.47255 \\ -0.08378 & -0.56403 & 1.64781 \end{pmatrix}$$

observe how we have large values on the diagonal and smaller but still significant negative values next to them - this is fairly typical.

The algorithm allows the colour saturation of the output image to be altered. It does this by transforming the colour matrix $C_{original}$ from above according to the following equation for the saturation parameter s.

$$C_{new} = A^{-1} \cdot \begin{pmatrix} 1 & 0 & 0 \\ 0 & s & 0 \\ 0 & 0 & s \end{pmatrix} \cdot A \cdot C_{original}$$

where the matrix A here converts from RGB to YCbCr, *viz.*

$$A = \begin{pmatrix} 0.299 & 0.587 & 0.114 \\ -0.169 & -0.331 & 0.500 \\ 0.500 & -0.419 & -0.081 \end{pmatrix}.$$

5.12.3. External API

The `Ccm` class is derived from `CcmAlgorithm` . This means it implements the following additional public methods.

CcmAlgorithm class API
`setSaturation(double saturation)` Sets a fixed saturation level, where 1.0 = normal saturation, 0.0 = greyscale and values greater than 1.0 imply increased colour saturation.

Code

```
Algorithm *algorithm = controller->GetAlgorithm("ccm");
CcmAlgorithm *ccm_algorithm = dynamic_cast<CcmAlgorithm *>(algorithm);
if (ccm_algorithm) {
    // Set colour saturation to 90% of the normal level.
    ccm_algorithm->setSaturation(0.9);
}
```

5.12.4. Parameters

Name	Default	Description
`ccms`	*Required*	List of colour temperatures followed by 3x3 colour matrix.
`saturation`	*Optional*	An optional piecewise linear function that allows colour saturation to vary with the estimated lux level of the image (for example, to reduce colour saturation in low light). It is specified by listing alternative lux and saturation values.

5.12.5. Metadata Dependencies

The CCM Control Algorithm requires

1. The `AwbStatus` metadata from the AWB algorithm, and
2. the `LuxStatus` metadata from the Lux algorithm.

The CCM Control Algorithm writes a `CcmStatus` object (`controller/ccm_status.h`) into the image metadata.

5.13. Sharpening

5.13.1. Name

`rpi.sharpen`

5.13.2. Overview

The Sharpening Control Algorithm gives rudimentary control over the Sharpening hardware block" in the ISP according to the following simple model.

The parameters are normally varied automatically according to the mode currently being used by the camera. The default configuration and values should normally be sufficient for many purposes, although they may need adjustment by hand for some sensors.

5.13.3. External API

The `Sharpen` class is derived from `SharpenAlgorithm` . This means it implements the following additional public methods.

SharpenAlgorithm class API
`setStrength(double strength)` Sets the sharpening strength, where 1.0 = normal sharpness, 0.0 = no sharpening at all and values greater than 1.0 imply increased sharpening.

Code

```
Algorithm *algorithm = controller->GetAlgorithm("sharpen");
SharpenAlgorithm *sharpen_algorithm = dynamic_cast<SharpenAlgorithm *>(algorithm);
if (sharpen_algorithm) {
    // Set sharpening to 80% of the normal level.
    sharpen_algorithm->setStrength(0.8);
}
```

5.13.4. Parameters

Name	Default	Description
`threshold`	1.0	Threshold below which pixels are not sharpened. Use values < 1 to sharpen more low contrast pixels; use values > 1 to sharpen fewer low contrast pixels.
`strength`	1.0	Slope of the rate at which the amount of sharpening applied increases. Use values < 1 to reduce the rate at which sharpening increases; use values > 1 to apply increasing amounts of sharpening more quickly.
`limit`	1.0	Limit to the total amount of sharpening applied. Use value < 1 for a lower maximum amount of sharpening; use values > 1 to allow larger amounts to sharpening to be applied.

5.14. Auto Focus

5.14.1. Name

`rpi.af`

Note: This algorithm is only available for cameras with an associated lens driver. To implement Continuous AF, it requires a PDAF sensor, such as Sony IMX708.

5.14.2. Overview

The AF algorithm is responsible for controlling a lens driver to achieve manual or auto focus. It has three modes of operation:

- In **Manual** mode, it simply sets the lens to focus at a distance specified by the user. Note that many lens drivers are not accurately calibrated for absolute distances (but are usually repeatable);
- In **Auto** mode it will search for the best focus when triggered by the user, using either Contrast Detection Auto Focus (CDAF) or Phase Detection Auto Focus (PDAF);
- In **Continuous** mode it implements a feedback loop to try to maintain the best focus - this mode works only when PDAF data are available;

5.14.3. External API

The `Af` class is derived from `AfAlgorithm` , and implements the following additional public methods:

AfAlgorithm class API
`setRange(AfRange range)` Sets the range of lens travel, which should be one of { `AfRangeNormal` , `AfRangeMacro` , `AfRangeFull` }. Default is `AfRangeNormal` .
`setSpeed(AfSpeed speed)` Sets the convergence speed, which should be one of { `AfSpeedNormal` , `AfSpeedFast` }. Default is `AfSpeedNormal` .
`setMetering(bool use_windows)` Controls whether user-defined windows are used for AF. When disabled, a default central window is used.
`setWindows(libcamera::Span<libcamera::Rectangle const> const & windows` Set AF windows. These are defined in pixels relative to the full sensor; i.e. the same coordinate system in which crop rectangles are defined. If no windows are specified or if all windows are outside the current mode's maximum crop region, a default central window is used.
`setMode(AfMode mode)` Set the focus mode, which should be one of { `AfModeManual` , `AfModeAuto` , `AfModeContinuous` }.
`setLensPosition(double dioptres, int32_t *hwpos)` For use in Manual mode: focus at the given position, specified as the inverse of a distance in metres. If `hwpos` is not null and if the function returns true, `*hwpos` will be set to the V4L2 control value to which the lens should immediately be driven.
`triggerScan()` For use in Auto mode: start a scan.
`cancelScan()` For use in Auto mode: cancel any incomplete scan.
`pause()` For use in Continuous mode: pause or resume continuous focus.

5.14.4. Parameters

Name	Default	Description
`ranges`		Parameters that depend on `AfRange` :
`normal`		Parameters applied in `AfRangeNormal` :
`min`	0.0	Inverse of furthest focus distance in metres.
`max`	12.0	Inverse of nearest focus distance in metres.
`default`	1.0	Default inverse distance.
`macro`	*Same as normal*	Parameters applied in `AfRangeMacro`
`full`	*Union of normal,macro*	Parameters applied in `AfRangeFull`
`speeds`		Parameters that depend on `AfSpeed` :
`normal`		Parameters applied in `AfSpeedNormal` :
`step_coarse`	1.0	Coarse step size in dioptres (inverse metres), for CDAF.
`step_fine`	0.25	Fine step size in dioptres (inverse metres), for CDAF.
`step_frames`	4	Number of frames to linger between steps, for CDAF, to allow the lens to settle and for latency.
`contrast_ratio`	0.75	Contrast threshold, relative to peak value, below which to terminate coarse scan and start fine scan.
`pdaf_gain`	−0.02	Feedback coefficient relating PDAF phase to lens movement per frame. The correct sign and magnitude will depend on how phase is encoded by the sensor. For a sensor without PDAF, this must be 0.
`pdaf_squelch`	0.125	Lens movement per frame below which to suppress small movements in Continuous mode, or to terminate a scan in Auto mode with PDAF. Ignored without PDAF.
`max_slew`	2.0	Maximum lens movement per frame.

Name	Default	Description
`pdaf_frames`	20	Number of frames to run the PDAF loop for each triggered scan in Auto mode. For a sensor without PDAF, this must be 0.
`dropout_frames`	6	Number of frames with low PDAF confidence before PDAF will be abandoned in favour of a CDAF-based scan. For a sensor without PDAF, this must be 0.
`fast`	*Same as normal*	Parameters applied in `AfSpeedFast`
`conf_epsilon`	8	PDAF confidence global threshold. The scale of this and the following two parameters will depend on how confidence is encoded by the sensor. Ignored without PDAF.
`conf_thresh`	16	Per-region confidence threshold. Regions with less than this PDAF confidence are ignored. Ignored without PDAF.
`conf_clip`	512	Per-region confidence upper limit. Regions with more than this PDAF confidence do not receive any additional weight. Ignored without PDAF.
`skip_frames`	5	Frames to skip at startup or after a mode switch, before phase and contrast data can be considered reliable.
`map`	[0.0, 445, 15.0, 925]	PWL function which relates inverse distance to hardware lens setting (`V4L2_CID_FOCUS_ABSOLUTE`). Also defines the minimum and maximum possible lens position. It must be specified for each (sensor,lens) combination.

5.14.5. Metadata Dependencies

The Auto Focus algorithm reads the `focusRegions` statistics, which are used for CDAF. Where available, it also reads the `pdaf.regions` image metadata, generated by the sensor via its camera helper.

The output of the Auto Focus algorithm is the `af.status` image metadata, which encodes its state and, optionally, the setting to which the lens should be driven.

Lens settings should be acted upon as soon as possible and should not wait until the associated frame is at the head of the queue for display.

5.15. Chromatic Aberration Correction

5.15.1. Name

`rpi.cac`

Note: This algorithm is only available on Pi 5.

5.15.2. Overview

This hardware block is able to perform a slight resampling of the red and blue pixels in the Bayer pattern in order to reduce lateral chromatic aberration. It is accomplished by shifting the red and blue pixels up to 2 pixels in any direction.

These shift values are defined on an 8x8 grid covering the image. The shift values for each pixel are interpolated bi-linearly from this table.

5.15.3. External API

The `Cac` class does not currently implement any extra methods.

5.15.4. Parameters

Name	Default	Description
`lut_rx`	*Optional*	Table of 8x8 pixel displacements, each value lying between −2.0 and 2.0. This table defines the amount of shift applied to red pixels in the x (horizontal) direction. When unspecified, all the values default to zero.
`lut_ry`	*Optional*	As `lut_rx` but for the y (vertical) direction.
`lut_bx`	*Optional*	As `lut_rx` but for the blue pixels.
`lut_by`	*Optional*	As `lut_ry` but for the blue pixels.
`strength`	1.0	A scale factor applied to the pixel displacements as they are loaded.

5.16. HDR (High Dynamic Range)

5.16.1. Name

`rpi.hdr`

5.16.2. Overview

The HDR algorithm can support applications that require images with a higher dynamic range than normal. It does so by running the AEC/AGC algorithm in a specially configured manner, and may then (on a Pi 5) merge and apply tonemapping to the results.

Normally AGC runs in a single channel configuration, metering for a compromise level of brightness for the image as a whole. However, we can define separate AGC channels to meter, for example, for the image highlights (because this will shorten the exposure time it is often referred to as the “short” channel), or for the darker parts of the image (similarly, the “long” channel).

The HDR algorithm can be set up to work with AGC running with a single different channel from the usual one (usually a “short” channel), or using a number of different channels which are run on alternating frames (such as a “short” and a “long” channel).

On a Pi 5 the images from these different channels can (optionally) be merged and tonemapped, whereas on a Pi 4 or earlier device, this feature is not available. In this case the application will receive alternating “short” and “long” exposure images, being, respectively, under- and over-exposed. It will be up to the application to decide how to deal with the differently exposed images.

Tuning the HDR algorithm therefore consists of two steps.

- The AGC channel or channels must be configured in the `rpi.agc` section of the tuning file.
- The HDR algorithm must be configured to advertise the AGC channels it wants to use, and on a Pi 5 it may additionally be configured how to merge and tonemap images.

The tuning files supplied for a Pi 5 contain AGC channel and HDR configurations that should be useful in a variety of situations.

AGC Channels

The AGC channels are as follows:

- Channel 0 - normal AEC/AGC metering.
- Channel 1 - HDR short channel metering. This channel will attempt to expose for image highlights, but not to the point of crushing the vast majority of the image to near darkness.
- Channel 2 - HDR long channel metering. This channel will meter for the less bright parts of the image, whilst not straying too far from the short channel (channel 1).
- Channel 3 - A “night” mode metering channel. The tuning file supplies a subtle HDR mode which modestly brightens up very dark images, whilst also recovering mildly over-exposed highlights.

HDR Modes

HDR configurations are each named and independent, and the ones defined to use these channels are:

- `off` - HDR is disabled. This HDR “mode” engages AGC channel 0 as the sole AGC channel.
- `SingleExposure` - HDR mode that accumulates multiple “short” exposure images using AGC channel 1. This is recommended on Pi 5 for most HDR applications.

- `MultiExposure` - HDR mode using short and long images (AGC channels 1 and 2). This mode can be used only where a scene is completely static, does not really work for video, and the `SingleExposure` mode should nearly always be preferred.
- `Night` - a “mild” HDR mode for low light imaging.

Finally there is also a special `MultiExposureUnmerged` mode which runs a short and long channel just like the `MultiExposure` mode, but applies no image merging and tonemapping (the application receives unmodified under- and over-exposued images). This mode is the only mode that could be used on a Pi 4 (or earlier device). It can also be used on a Pi 5 where an application wants to take control of the HDR processing.

5.16.3. External API

The `Hdr` class is derived from `HdrAlgorithm` . This means it implements the following additional public methods.

HdrAlgorithm class API
`setMode(std::string const &modeName)` Sets the HDR mode to one of the ones named in the tuning file, such as `Off` , `SingleExposure` and so forth.
`getChannels()` Get the list of AGC channels for this mode. The returned list should be passed to the AGC algorithm’s `setActiveChannels` method.

Code

```
Algorithm *algorithm = controller->GetAlgorithm("sharpen");
SharpenAlgorithm *sharpen_algorithm = dynamic_cast<SharpenAlgorithm *>(algorithm);
algorithm = controller->GetAlgorithm("agc");
AgcAlgorithm *agc_algorithm = dynamic_cast<AgcAlgorithm *>(algorithm);
if (hdr_algorithm && agc_algorithm) {
    hdr->setMode("SingleExposure");
    agc->setActiveChannels(hdr->getChannels());
}
```

5.16.4. Parameters

Name	Default	Description
`cadence`	*Required*	The sequence in which to run AGC channels on the camera frames. For example, `{ 1, 2 }` will run AGC channel 1 and then channel 2 on alternate frames.
`channel_map`	*Required*	A table mapping the HDR terms “short” and “long” onto the AGC channels used in the `cadence` .
Note:		all the subsequent fields apply only to Pi 5
`stitch_enable`	0	Whether to stitch multi-channel images together. This should be set to 1 when there are two (“short” and “long”) AGC channels, and otherwise zero. Though as noted, the single (and not multi) exposure mode of HDR operation is recommended.
`tonemap_enable`	0	Whether to apply tonemapping. Setting this to 0 will make a Pi 5 run like an earlier device, where no tonemapping occurs and the frames for the different AGC channels are returned directly.
`spatial_gain`	2.0	The darkest areas of the image will have this gain applied to them. By default, areas where the brightness level is between 0 and 1% of the full range will have this gain applied, tapering down to a gain of 1.0 at 6% of the full range.
`spatial_gain_curve`	*Optional*	A more detailed way of setting the spatial gain. Here the entire piecewise linear function relating brightness to gain can be specified. For example, `{ 0, 2.0, 0.01, 2.0, 0.06, 1.0, 1.0, 1.0 }` would be equivalent to setting `spatial_gain` to 2.
`diffusion`	3	The spatial gain is defined on a grid with granularity 32x32 across the image. This value determines across how many of these individual cells the gain can "diffuse". Too low a value leads to obvious halos, whereas too large a value reduces the overall usefulness of the spatial gain.

Name	Default	Description
`hi_quantile_targets`	*Optional*	A list of *quantile*, *target* pairs of numbers. When a scene does not show much saturation, as judged by the quantile values in this list not reaching the targets, then the tonemapping tries to use gain in place of tonemapping curves, as this better preserves local contrast. The default value for this field is `{ 0.95, 0.65, 0.5, 0.28, 0.3, 0.25 }` , meaning that the 95% point of the histogram should reach the 65% part of the dynamic range, and so on.
`hi_quantile_max_gain`	1.6	The maximum gain that can be applied as a result of the “high quantile” analysis.
`quantile_targets`	*Optional*	Somewhat like the `hi_quantile_targets` , these are targets for quantiles at the low end of the dynamic range that the algorithm will attempt to achieve through a power law tonemap, rather than through gain. It defaults to `{ 0.2, 0.03, 1.0, 0.15 }` , meaning that the 20% quantile should ideally reach at least 3% up the dynamic range, and so on.
`power_min`	0.65	The minimum allowed value of the exponent calculated by the low quantile analysis. Decreasing this value would allow stronger tonemapping to be used.
`power_max`	1.0	The maximum allowed value of the exponent calculated by the low quantile analysis.
`contrast_adjustments`	`{ 0.5, 0.75 }`	The power law used for tonemapping tends to create slightly washed out images because of the steep gains at the bottom. These factors multiply the output values at the bottom end of the tone curve, so as to lower them. The values can be changed, or more values added, so as to apply further up the curve.

5.16.5. Metadata Dependencies

The HDR algorithm requires the `agc.status` metadata and `alsc.status` metadata to be available.

The HDR Control Algorithm writes a `TonemapStatus` object (`controller/tonemap_status.h`) into the image metadata when tonemapping is enabled, and a `StitchStatus` object (`controller/stitch_status.h`) when stitch is enabled.

5.17. Software Camera Synchronisation

5.17.1. Name

5.17.2. Overview

Raspberry Pi’s libcamera implentation has the ability to synchronise frames between multiple cameras purely through software. That is, the start times of the frames will be brought as close together as possible.

The scheme works by having one camera designated as the *server* which broadcasts timing messages onto the network. Other cameras - which may be attached to the same Pi or a different one - are known as *clients* and listen out for server messages. The clients will then adjust their camera’s frame timing so as to bring it, as far as possible, into lock step with the server. There is no back-channel, so the server has no knowledge of the state or even existence of any clients.

Once it is running, the server waits for a fixed number of frames before assuming that clients have had time to synchronise - so it’s usually best to start clients before the server. All devices will then signal the camera application that is running that the moment at which frames are guaranteed to be synchronised (the “synchronisation point”) has happened.

Name	Default	Description
`group`	"239.255.255.250"	For the server, the IP address to which timing messages are broadcast. For clients, the address they listen to for timing messages.
`port`	10000	The port number to go with the `group` IP address.
`sync_period`	30	The number of frames the server waits before sending out another timing message. This parameter is ignored by clients.

Name	Default	Description
`ready_frame`	100	The number of frames the server waits before declaring that all devices are synchronised. This parameter is also ignored by clients (as they learn it from the server messages).
`min_adjustment`	50	The time, in microseconds, below which clients will not attempt to adjust frame durations to match the server. This parameter is ignored by servers.

5.17.3. Metadata Dependencies

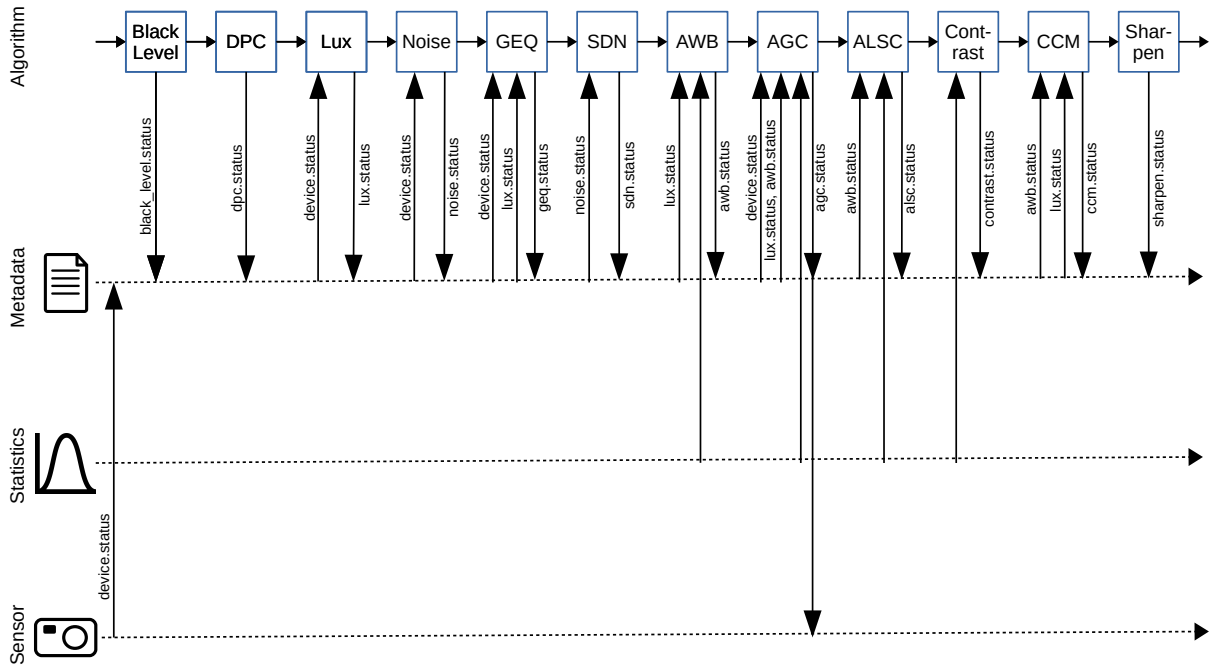
The Sync control algorithm expects image metadata to contain a `SyncParams` data structure (`controller/sync_status.h`) under the key `sync.params` which contains a wallclock timestamp, measured in microseconds, corresponding to when the most recent frame started.

The Sync control algorithm writes a `SyncStatus` object (`controller/sync_status.h`) into the image metadata. This contains information about whether the sychronisation point has been reached, how much frame length adjustment to apply, and so forth.

5.18. Metadata and Statistics Usage

We finish with a diagram (Figure 16) showing how image metadata is both read and written to by the various stages of procesing.

Figure 16.
Usage of image metadata and other information by ISP control algorithms on a Pi 4. Pi 5 operates mostly in an identical manner.



In Figure 16, observe that:

- Device metadata (`device.status`) is always supplied from the camera information.
- After AGC/AEC has run, updated exposure and gain values are fed back to the camera (though this happens outside of the algorithm itself).
- Some algorithms use image statistics but many do not. Many use only the `device.status` from the image metadata.
- All algorithms write `status` information back to the image metadata.

6. Camera Tuning Tool

6.1. Overview

The Camera Tuning Tool (CTT) is a Python program designed to produce a fully working camera tuning JSON file from a relatively small set of calibration images. Once the tool has run, there should be either no or only minimal further tweaking to the JSON file required in order to obtain the desired image quality. The tuning algorithms are furthermore designed to work with a minimum amount of expensive or specialised equipment. The processes required in creating a finished camera tuning are as follows.

1. Firstly, a functional V4L2 camera driver must be written (see Chapter 3). For the purposes of writing the camera driver, an uncalibrated tuning file can be copied which should provide recognisable images. Copy one of `vc4/data/uncalibrated.json` (Pi 4 and earlier devices) or `pisp/data/uncalibrated.json` (Pi 5), depending on the platform you are using. Care should also be taken that the black level listed in the file is adjusted to match the black level specified in the sensor data sheet (and scaled up to a 16-bit range).
2. The set of calibration images must be captured. Again, this should use the uncalibrated tuning file. There are two types of calibration images, those with a Macbeth chart, and a further set of completely uniform images for measuring lens shading.
3. On a Pi 5 you may optionally capture some images to tune the Chromatic Aberration Correction (CAC) block. These images can also be omitted. This feature does not exist on Pi 4 or earlier devices.
4. With the calibration images all correctly named and stored in a folder, the CTT can be run. The CTT finds Macbeth charts in images automatically and uses them to measure noise profiles, green imbalance, white balance and colour matrices.
5. The output JSON file of the CTT can be used directly, possibly with minor further tweaking.

6.2. Raspberry Pi *rpicam-apps*

Image capture is most conveniently performed using Raspberry Pi's *rpicam-apps*. For readers who are familiar with the existing *raspistill* and *raspivid* applications, the *rpicam-apps*, including *rpicam-still* and *rpicam-vid* are *libcamera*-based replacements. Though not completely identical in functionality, they aim to be very similar.

The *rpicam-apps* should be installed by default on any Raspberry Pi OS. For more information about them, please refer to https://www.raspberrypi.com/documentation/computers/camera_software.html.

6.3. Software Installation

The CTT runs under Python 3.11, and will install all required dependencies in a virtual environment:

```
Code
sudo apt install python3-pip
python3 -m venv venv
. venv/bin/activate
pip3 install rpi-ctt
```

6.4. Equipment

The additional equipment required for camera tuning is described below.

6.4.1. X-rite (Macbeth) Colour Checker

These charts are very well known, and one is shown i. They normally cost less than £100.

Figure 17.
X-rite Macbeth Chart.



6.4.2. Colorimeter

Some kind of colorimeter is required in order to measure the colour temperature and light level of the calibration images. One such is shown in Figure 18, costing normally a few hundred pounds.

Figure 18.
A Colorimeter.



Depending on the application it may be possible to get by with significantly simpler instruments. For example, a mobile phone may well report colour temperature and light level, and while its readings may be quite approximate, it may suffice for certain use cases.

6.4.3. Integrating Spheres and Flat Field LEDs

There are some relatively sophisticated pieces of equipment - integrating spheres and flat field LED lamps - that can help calibrate lens shading, and these are likely to prove beneficial when they are available. The CTT is designed to average multiple images for each colour temperature, so quite reasonable results can be obtained with a cheap LED video light costing less than £50, possibly with an extra diffuser. It's convenient if the lamp can generate a couple of colour temperatures, approximating to "indoor" and "outdoor".

6.5. Capturing Calibration Images with *libcamera*

The tuning process requires raw images to be captured and fed into the tuning tool. We recommend Raspberry Pi's *rpicam-apps* for this purpose, as described earlier, however the *qcam* application supplied with *libcamera* itself is also suitable. Both applications will capture the required full resolution raw files (in DNG, or Adobe "digital negative" format).

6.5.1. Capturing a Raw Image

Camera tuning files are kept in the `vc4/data` (Pi 4 or earlier devices) or `pisp/data` (Pi 5) folder under the Raspberry Pi IPA folder, and the control algorithms expect to find a JSON file there, named according to the sensor name exposed in the driver, for every

camera that is used. When tuning a new camera, of course, we don't have one so we provide a tuning file for *uncalibrated* cameras which can be copied. For example, when tuning a new sensor named `xyz123` we simply copy `uncalibrated.json` to `xyz123.json`, all in the same folder.

The newly copied uncalibrated tuning will produce recognisable images from pretty much any sensor, but there is just one single parameter that *does* have to be configured before use. This is the *black level*, and it must be set to the correct value from the sensor datasheet scaled up to 16 bits. So if your sensor produces 10-bit samples, and the black level is 64, then the correct value in the new uncalibrated tuning file (`xyz123.json`) will be 4096.

To capture a DNG using *rpicam-still*, enter (substitute an image name of your choice for `image.jpg`)

```
Code
rpicam-still -r -o image.jpg
```

which will produce a full resolution jpeg file named `image.jpg` and a DNG file `image.dng` (the same name but with the extension `.dng`). This second file contains the raw data that produced the jpeg.

Note that the CTT may sometimes discard an image if it is too over-exposed to be useful, and it will report this in the console log. Mostly the tuning produced should still be reasonable, but you may wish to retake those particular images with a lower exposure, for example:

```
Code
rpicam-still -r -o image.jpg --ev -1.0
```

This will reduce the exposure by one stop.

6.6. Image Capture Requirements

As we have already explained, there are two (on a Pi 5, optionally three) distinct types of calibration images. Because we are most likely to be using the uncalibrated tuning, it is possible that the colour balance may appear incorrect in some of the images. Rest assured that this makes no difference - it is only the raw camera data (which is unaffected) that is used by the tuning tool; the image displayed while capturing is purely for convenience and to aid framing each shot.

6.6.1. Macbeth Chart Images

The aim here is to capture images containing a Macbeth chart. The CTT locates the Macbeth charts automatically, therefore the chart should be

- Reasonably central,
- Not too small,
- Reasonably straight, and
- Unobscured, except perhaps a little bit at the edges (where maybe it is being held).

The images should be captured under a range of different colour temperates spanning at least the operating conditions the camera is expected to encounter. For a camera system to be used across a wide range of conditions, this might include lamps from about 2500K up to about 8000K. It is acceptable to capture images both in a camera lab that can generate the different illuminants, and/or out in real world operating conditions. In general, the closer the calibration pictures to the final operating environment, the better the results are likely to be. Below are a couple of acceptable calibration pictures.

Figure 19.
Acceptable Macbeth Chart calibration images.





The left hand image is taken in a small purpose-built camera lab/studio, whist the second was merely taken out in an office lobby. We would probably not recommend going very much smaller than the chart in this left hand image. Observe further that the right hand image shows some barrel distortion - at the level shown the CTT is still able to work with the image.

Once captured, images should be named so as to contain the colour temperature (followed by 'k' or 'K') and the lux level (followed by 'l or 'L') in the filename. Valid filenames might include

- `imx219_2954k_1749l.dng` for an image captured in 1749 lux with a colour temperature of 2954K, or
- `1749L_2954K.dng` for the same image.

6.6.2. Lens Shading Images

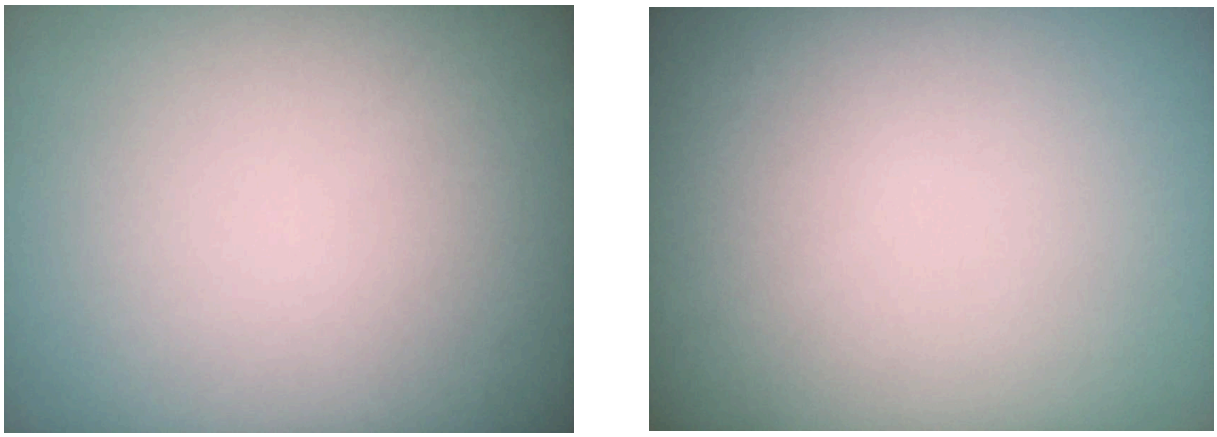
As with the Macbeth calibration images, images should be captured in different colour temperatures depending on what illuminants are available (though not necessarily the same ones). Images that accurately reflect and cover the range of operating conditions are likely to be beneficial. Here, however, the image must be of a completely featureless, flat and uniform surface, such as an LED video light. Again, the filenames should contain the colour temperature (followed by 'k' or 'K'). The lux level does not need to be recorded on this occasion, though the filenames **must** include `alsc` so that these images can be distinguished from the Macbeth chart images as they are loaded.

Where specialised equipment is available a single lens shading image for each colour temperature may be sufficient. Otherwise images should be captured where the luninance variation across the scene is as low as practicably possible, and the CTT will expect a number of calibration images for each colour temperature, where each image should be taken with the camera rotated or positioned differently. For example, for each colour temperature we might submit

- 2 images each, one where the camera is the usual way up and the other where the camera is held upside down, or
- 4 images each, where each image positions the camera distinctly on an LED video light.

The CTT averages the images for each colour temperature, thereby reducing the effect of the uneven illumination on the scene and making the intrinsic lens-produced colour and luminance shifts more discernible.

Figure 20.
Lens shading images `alsc_3850k_1.dng` and `alsc_3850k_2.dng`. In this case the images differ slightly because they have been placed differently on an LED light for capture.



The single luminance correction table required by the ALSC algorithm is calculated by averaging together all the lens shading calibration images, for all colour temperatures.

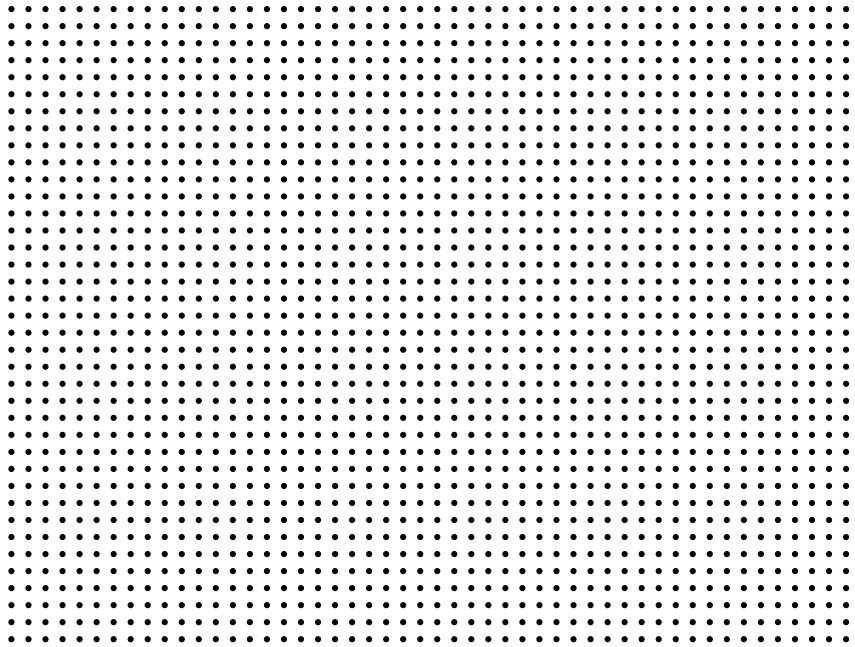
6.6.3. Chromatic Aberration Correction (CAC) Images

This feature is only available on Pi 5 and can be ignored on other platforms. Even on Pi 5, it is optional to supply these images. If none is supplied, the CAC function in the ISP will be left disabled by the final tuning file.

One or more CAC images can be supplied, and they will all be identified by having the string `cac` in the file name. Where more than one image is supplied, they are averaged together to form the final CAC adjustment tables. The CAC correction feature only works for lateral chromatic aberration.

The images should be of the `cac_chart.svg` file, which should be printed out with a reasonable quality printer that does not create any colour fringes. Normally best results will be obtained by printing the image as large as possible so that the camera is further away, preferably at least A3 size. We also recommend taking care to adjust the focus of the lens where this is variable.

Figure 21.
This image, `cac_chart.svg`, should be supplied with the Camera Tuning Tool.



6.7. Creating the Tuning

6.7.1. Collecting the Files

Once all of the calibration images have been captured, they should be placed together into a single folder with no other image files. The directory listing below shows an example of what such a calibration folder might contain.

```
Code
2498K_106l.dng
2811K_403l.dng
2911K_1208l.dng
2919K_605l.dng
3627K_1247l.dng
4600K_998l.dng
5716K_1069l.dng
8575K_170l.dng
alsc_3000K_1.dng
alsc_3000K_1u.dng
alsc_3850K_1.dng
alsc_3850K_1u.dng
alsc_6000K_1.dng
alsc_6000K_1u.dng
```

In this example we have eight Macbeth chart images covering eight different illuminants and three pairs of lens shading images covering only three colour temperatures, and where the second of the pair is rotated 180 degrees compared to the first. Please be aware that there should be only a single Macbeth chart image at each distinct colour temperature.

6.7.2. Running the Tool

The tuning tool can be found under the root *libcamera* directory in `utils/raspberrypi/ctt` and can be run by executing `ctt.py` found there. It takes the following arguments

Name	Default	Description
`-i`	*Required*	Path of the folder containing the calibration images.
`-o`	*Required*	Name of the output JSON file to contain the camera tuning.
`-t`	*vc4*	The target platform for the tuning file, either `vc4` (Pi 4 or earlier devices) or `pisp` (Pi 5).
`-c`	*Optional*	The CTT takes a JSON file containing its own configuration options (see below). If omitted, default parameters are used.
`-l`	`ctt_log.txt`	Name of a log file to which diagnostic output from the tool is written.

For example, one might, in the *libcamera* folder `utils/raspberrypi/ctt`, enter

```
Code
./ctt.py -t vc4 -i ~/imx219_calibration_folder -o imx219.json
```

The tool takes quite a few seconds to run, reporting its progress during this time. The default configuration should normally be appropriate, but an example CTT configuration file is included below (also `ctt_config_example.json` in the same folder) for completeness.

```
Code
{
  "disable": [],
  "plot": [],
  "alsc": {
    "do_alsc_colour": 1,
    "luminance_strength": 0.5
  },
  "awb": {
    "greyworld": 0
  },
```

```
  "blacklevel": -1,
  "macbeth": {
    "small": 0,
    "show": 0
}
```

If necessary the fields here can be modified, as follows.

Name	Description
`disable`	Do not tune the listed stages. Generally should be left as `[]` .
`plot`	Show plots during the tuning process.
`do_alsc_colour`	Whether to calibrate for colour shading correction. Would be omitted for a monochrome sensor.
`luminance_strength`	Luminance shading strength to specify in the final output file, between zero and one (inclusive).
`greyworld`	Whether to request Grey World AWB in the tuning file, or the (usual) Bayesian algorithm.
`blacklevel`	Black level to use for this sensor (out of 16 bits). The value −1 indicates to attempts to deduce it from the input file, if possible. Otherwise it is recommended to enter the correct black level obtained from the sensor data sheet.
`small`	Whether to search for extra small Macbeth charts in images (takes slightly longer).
`show`	Whether to show locations of Macbeth charts found in images.

6.8. Tuning only the Lens Shading Correction

Sometimes it's necessary to use a different lens with a camera module from the one that was originally supplied, and for which the system has been calibrated. In these circumstances, you may find that the Lens Shading Correction (LSC) tables need to be re-tuned.

This proceeds in a very similar fashion to the full camera tuning just described, with the exception that you *only* need to collect the lens shading calibration images (with names starting "alsc"). So you will need some kind of flat field target (such as an LED video lamp), but you won't need a colour checker chart.

Once you have placed these images into a folder you need to run

```
Code
./alsc_only.py -t pisp -i /path/to/alsc-images -o alsc.json
```

(or use `vc4` instead of `pisp` for platforms earlier than the Pi 5).

Finally, you should locate the tuning file being used by your system (usually under `/usr/share/libcamera/ipa/rpi` , and be sure to select the correct target `pisp` or `vc4`) and replace the existing `"rpi.alsc"` parameter block by the equivalent one from the new JSON file (`alsc.json` in our example). Do take care to preserve the same levels of curly braces everywhere as compared to the original file.

If you find yourself editing system files like this, we would always recommend taking a backup first as your camera may fail to start if the tuning file has inadvertently become invalid.

6.9. Tweaking the Tuning produced by the Tool

6.9.1. Blocks not Tuned

There are a handful of blocks that the CTT does not actually tune but merely outputs default parameters. In these cases the default parameters should normally be sufficient (they are already "adaptive" to the image in some sense) and there would be no particular reason to change them other than as part of any more general aesthetic alterations to the tuning file. These blocks are:

1. Defective Pixels (`rpi.dpc`) - this will be left at "normal" correction strength which should usually be appropriate.
2. Spatial Denoise (`rpi.sdn` or the `sdn` block" in `rpi.denoise`) - there is no particular tuning for this block", however, as it inherits and uses the image noise profile which *is* tuned, there should be no particular need to change anything here.
3. Contast (`rpi.contrast`) - the CTT outputs a gamma curve though in fact it is merely a fixed "sensible" gamma curve. No actual measurements are made in order to deduce it, though again, it should work well in nearly all circumstances.

4. Sharpening (`rpi.sharpen`) - the default parameters have been set to give reasonable results with our existing sensors. However, it is possible that other sensors may require some manual tuning. For example, for lower resolution sensors with larger pixels we might expect to have to reduce the amount of sharpening.
5. Auto Focus(`rpi.af`) - by default this algorithm is not enabled. If enabled by hand, note that the default parameters are appropriate for a Raspberry Pi v3 camera and may give poor results with other cameras. Step sizes and delays can safely be left at the defaults, at least for a VCM-type focus mechanism, but it is critical to tune the `map` between inverse distance and lens settings and, where PDAF is available, gain and confidence values.

6.9.2. Guidance on how to Tweak the Tuning

The CTT makes every effort to produce a JSON tuning file that can be used directly and with minimal further effort. Of course there occasions when we may wish to tweak it, either for purely aesthetic reasons (“I want less sharpening”, or “I want brighter colours”) or because the camera is used in a slightly different situation than the one tuned for (maybe it is used under an illuminant that was not in the calibration set, meaning the colours could be slightly “off”). We include some general guidance for each block“ on how it might be tweaked.

6.9.3. Black Level (`rpi.black_level`)

There is nothing really to tweak here as the black level should match the one in camera data sheet (scaled up to 16 bits).

6.9.4. Defective Pixel Correction (`rpi.dpc`)

There is only a single `strength` parameter here, which defaults to 1 (“normal”). If defective pixels are not apparent it may be worth setting this down to zero as this would in theory allow very slightly more detail to be recovered, although the effect would be very marginal.

If defective pixels *are* a problem then the value should be increased to 2. Note that testing should always include low light images as “weak” pixels tend to come crawling out of the woodwork when the analogue gain is high.

6.9.5. Lux (`rpi.lux`)

There should not be anything to tweak here. The values are derived from one of the Macbeth Chart calibration images and unless these were categorised with an incorrect lux reading it is hard to envisage what one might need to change.

6.9.6. Noise (`rpi.noise`)

The noise profile itself - defined by the parameters here - should generally not be changed. If more or less spatial denoise is wanted then it would be better to alter the Spatial Denoise parameters (`rpi.sdn` or `rpi.denoise`).

6.9.7. Green Equalisation (`rpi.geq`)

Test images will have to be examined to determine if there is a green imbalance problem, as evidenced by the “maze” artifacts discussed previously. If the calibration images make up a wider set of conditions than those in which the camera will be used, then it is possible the numbers here could be reduced. However, if maze artifacts are visible then they will need to be increased. The values here (offset and slope) define a straight line - if maze artifacts are seen in the darker parts of an image then increasing the offset may be more effective; if they are seen in the bright parts of an image then increasing the slope may be recommended.

Maze artifacts may depend on both the light level and colour temperature of an illuminant and it is hard to predict exactly when they are likely to be worst. Therefore a reasonably broad test set would be advised to verify any changes.

6.9.8. Spatial Denoise (`rpi.sdn` or `rpi.denoise`)

The Spatial Denoise control algorithm takes the previously calculated Noise Profile and derives ISP parameters from it. It has two parameters

- `deviation` which multiplies the standard deviation of the

pixel noise, and `strength` which determines how much of the original (noisy) pixel to blend back in. We can provide the following rules of thumb.

1. If an image is showing speckles of noise (the “noise speckles” we saw earlier) then the `deviation` (default value 3.2) should be inceased. Once there are no noise speckles there is no benefit in increasing this further (you will just “eat” image detail).
2. It may sometimes be possible to reduce the deviation, though careful testing would be required to ascertain whether the level is still sufficient.
3. Once the deviation is set appropriately, the amount of denoise is controlled with the `strength` (where zero effectively disables denoise altogether and one denotes maximum denoise). Increasing this will use more of the denoised pixel, though there may come a point where the image starts to look “plasticky” or “like an oil painting”. In this case the strength should be reduced, allowing a little of the original noise back.

Figure 22.

Too little denoise (left), and on the right the leaves look even more plasticky than they really are!



6.9.9. AWB (`rpi.awb`)

Problems related to colours are very difficult to debug without capturing an image, in the problematic conditions, with a Macbeth Chart

- so doing this should really be the first step. Once such an image is

obtained it may simply be worth re-running the CTT with the new image.

Secondly, the AWB algorithm has some global “sensitivity” parameters. These are primarily designed for coping with module to module variation among the cameras, but can also be used to shift the AWB tuning slightly towards either purple (increase `sensitivity_r` and `sensitivity_b` by 5 or 10 percent) or towards green (reduce them by a similar proportion). This accounts quite conveniently for users who may prefer a very slightly greeny-yellow tone to their colours and those who do not.

Figure 23.
Effect of setting sensitivities to 0.8 (left), 1.0 (centre) and 1.2 (right).



The other global parameters that we may wish to adjust based on personal taste are `whitepoint_r` and `whitepoint_b` . For instance, people often quite like a slightly “warmer” look to their images, so this may sometimes be a worthwhile change to make globally. The values can be experimented with, but setting `whitepoint_r` to 0.05 and `whitepoint_b` to -0.05 should produce a slight but noticeable warming.

Beyond this we have to start delving more into the depths of the algorithm. Questions to ask are:

- Does the illuminant of my problem picture actually lie on the CT curve? If not you might need to move at least part of it.
- Is it searching enough of the CT curve to find my illuminant?
- Is it that some of my illuminants are on the curve but others lie too far to the side? If both need to work you might need to increase the `transverse_pos` or `transverse_neg` values, which determine how far off the curve an illuminant may lie.

Sometimes it may simply be that the illuminant is in the feasible set but the image content is misleading the algorithm (it has a colour cast that can be interpreted as being caused by a different illuminant). In some cases the prior illuminant distributions can be tweaked to force the algorithm to a particular solution - for example, in bright conditions we bias the results quite strongly towards sunlight, whatever the image content. But bear in mind that AWB is fundamentally an under-constrained problem - you can never be totally sure that the white wall under a yellowish illuminant wasn’t really a yellow wall under a white-ish illuminant.

6.9.10. AGC/AEC (`rpi.agc`)

The CTT just outputs a fixed set of AGC/AEC parameters, though in practice this set works pretty well with any camera.

If the tuning is found to be globally a bit too dark or a bit too bright, this can easily be adjusted using the `base_ev` parameter. Setting this to a value less than one makes everything a bit darker whilst leaving everything else untouched, and a value greater than one makes everything a bit brighter.

Figure 24.
`base_ev` *at the default value 1.0 on the left, and at 1.414 - about half a stop - on the right.*



If more specific changes are required it is largely a matter of taste how to set:

- The weightings of the AGC regions (in the metering modes);
- The manner in which analogue gain and exposure are divided up in the exposure modes (possibly so as to limit the camera framerate);

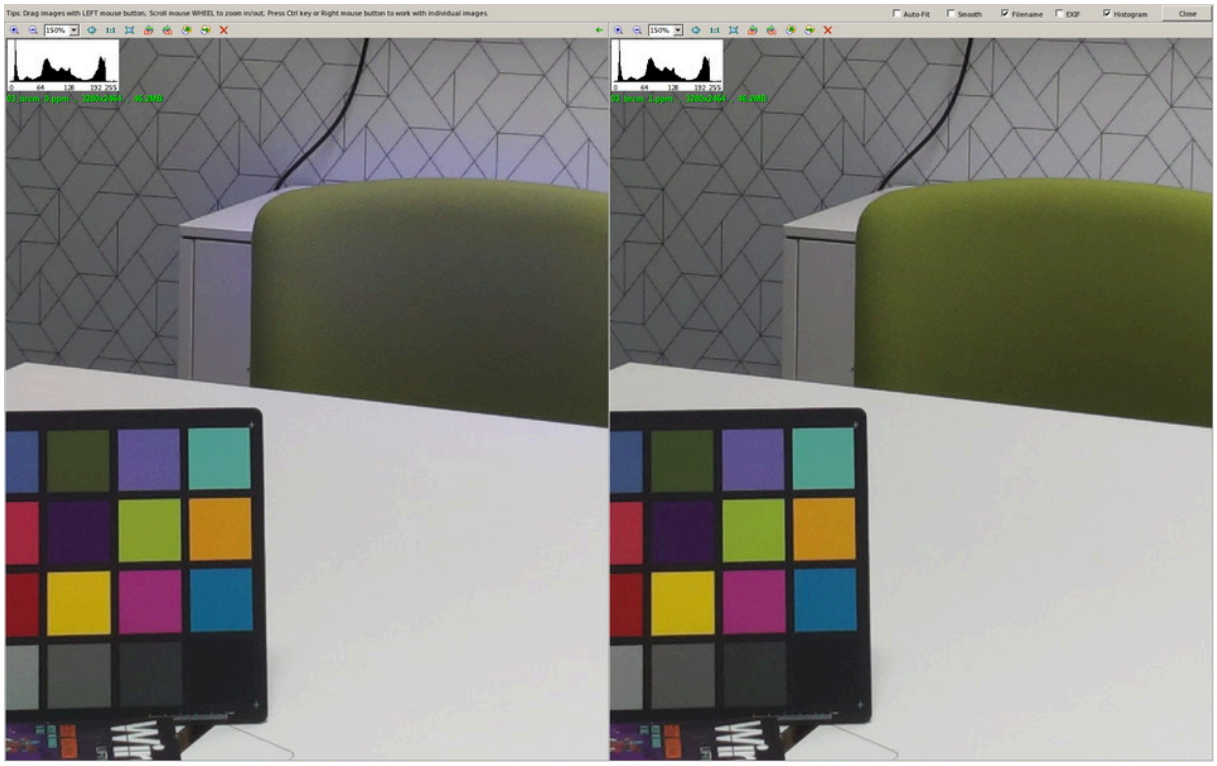
- The various Y target functions (which allow the Y target to change with lux level); and
- The histogram constraints (though an understanding of the section on the AGC algorithm will be required).

6.9.11. Automatic Lens Shading Correction (`rpi.alsc`)

There are a number of large tables here generated by the CTT, and it's unlikely that one would ever want to make changes to these by hand. However, there are a number of parameters that one might reasonably want to fiddle with.

- `omega` controls the amount of over-relaxation in the Gauss-Seidel method. The default value is 1.3; any value close to 1 is likely to work fine.
- `luminance_strength` determines the amount of luminance correction (recall that colour shading correction is always applied fully). Full luminance correction would be applied by the value 1.0, though in practice small values in the 0.5 to 0.8 range will look better.
- `n_iter` gives the number of Gauss-Seidel iterations. Generally the iterations stop before this limit anyway, so changing it does not have a great effect. The value zero can be useful as a way of turning off the adaptive algorithm completely (only the pre-measured calibrated tables will be applied).
- Note also that the pre-measured calibrations can be disabled either by deleting them from the file, or renaming `calibrations_Cr` to something like `x.calibrations_Cr` , and similarly for `calibrations_Cb` .
- The sigma values can be increased to force the adaptive algorithm to "mop up" more residual colour shading. Once you get to values beyond 0.01 there is a risk that some genuine colour variation in the image might start to get slightly smeared; such large values are therefore not recommended.

Figure 25.
`sigma` has the extreme value 0.03 on the left - spot the purple halo round the green chair.



6.9.12. Contrast (`rpi.contrast`)

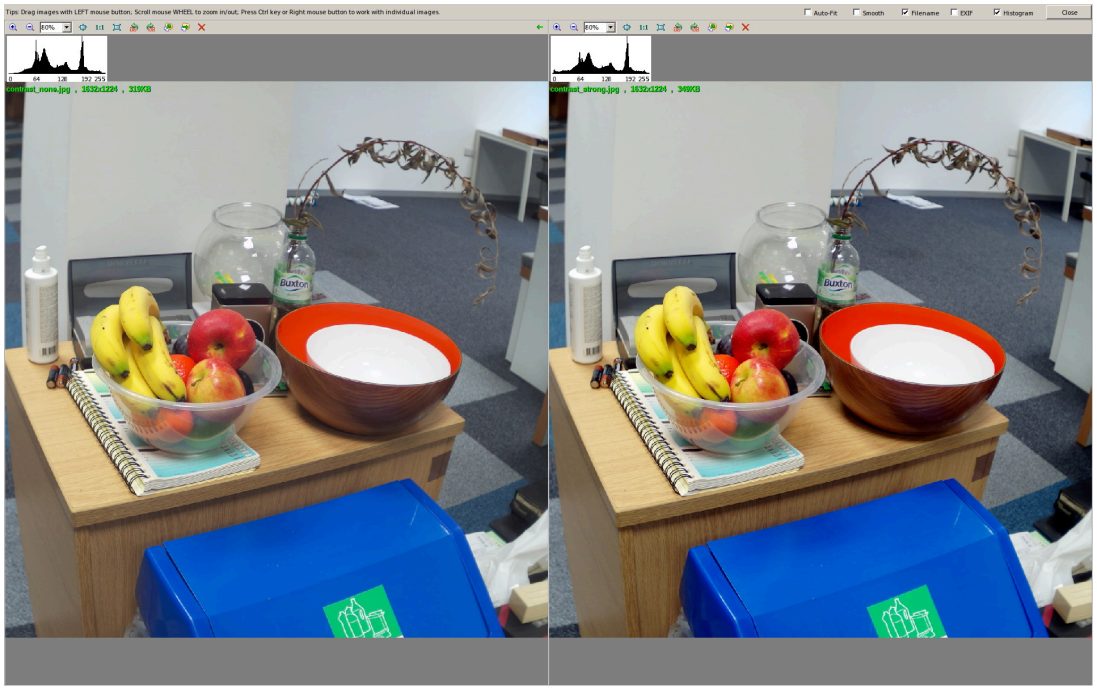
Firstly, the gamma curve can easily be changed or replaced by another curve, there is no constraint on how many points there must be or at what intervals they lie - any piecewise linear function will do. The values must of course occur in pairs (each pair of numbers is an (x,y) point), and obey the 16-bit dynamic range.

Setting `ce_enable` to zero completely disables any runtime adaptation of the gamma curve. When set to one the global contrast of an image can be increased by one or more of:

- Choosing a higher quantile value for `lo_histogram` (default 0.01).
- Choosing a lower level in the output range `lo_level` (default 0.015).
- `lo_max` (default 500) might need to be increased if the desired adjustment is larger than this value.

The high end of the dynamic range can be manipulated similarly using the “hi” (rather than “lo”) versions of these parameters, however, the effect is not usually as significant.

Figure 26.
No contrast enhancement (left) and strong contrast enhancement (right) - note the change to the image histogram.



6.9.13. Colour Correction Matrices (`rpi.ccm`)

Matrices are calculated by the CTT from the Macbeth Chart images. It fits them in RGB space as this is most simple, however, there would be other arguably better ways to accomplish this if one wished - for example measuring the fit error in Lab* space.

Otherwise there are some easy manipulations enabled by the `saturation` piecewise linear function. For instance one could globally reduce or increase the colour saturation. Adding the following would increase colour saturation globally by 10%.

Code

```
"saturation": [0, 1.1, 10000, 1.1]
```

6.9.14. Sharpening (`rpi.sharpen`)

All the parameters (`threshold` , `strength` and `limit`) assume default values of 1.0. If it is clear that significantly less (or more) sharpening is required, for example if the sensor is exhibiting rather different behaviour from that which we have been used to, it may be worth applying a few rounds of:

- *halving* all the parameters (if there was far too much sharpening), or
- *doubling* all the parameters (when there was far too little sharpening)

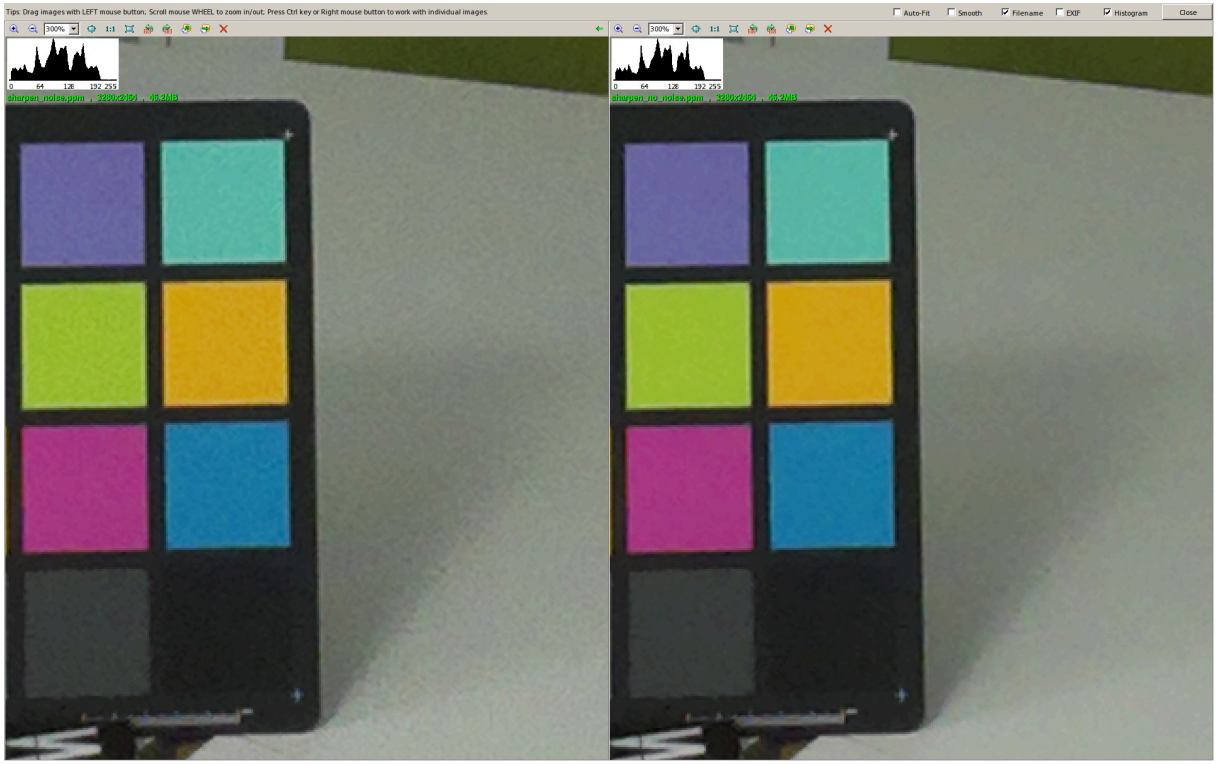
until the degree of sharpening appears to be approximately of the right order of magnitude. Once this has been achieved the parameters can be fine tuned. If more sharpening is required:

1. Consider lowering the `threshold` as this will cause more low-contrast detail to be picked up and sharpened.
2. Consider increasing `limit` - this will cause the strongest most obvious edges to be sharpened more.
3. Consider increasing `strength` - this will cause all details to be sharpened more up to the given `limit` .

In order to reduce the amount of sharpening, the reverse steps should be taken. In particular, if it looks as though noise is being picked up and sharpened unwantedly, then the `threshold` value needs to be raised.

Figure 27.

On the left too much residual noise is being sharpened because `threshold` was lowered.





Raspberry Pi

Raspberry Pi is a trademark of Raspberry Pi Ltd

Raspberry Pi Ltd