Cooke /i Technology

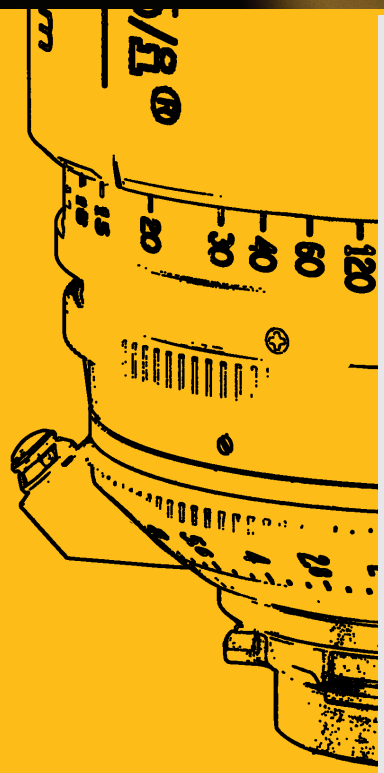

Cooke /i Lens Models Part III

Inertial, Shading and Spherical Distortion Data

2021

cookeoptics.com

Contents

Inertial Sensor Basic Data Processing

1 Introduction – Inertial Sensor Basic Data Processing

Cooke lenses with /i² and /i³ Technology contain a set of inertial sensors, which provide information about the 3D orientation and motion of the lens. This system has been designed primarily to support VFX applications, particularly as an aid to feature-based 3D camera tracking workflows.

This application demands high accuracy, and Cooke's system has been designed with this in mind. Each lens is factory calibrated to minimise the impact of manufacturing variation in the inertial sensors, data processing avoids introducing artefacts into the stream, and timing data supports synchronisation between inertial and video data. This design provides the best possible inputs for camera tracking, but it also means that lens outputs are essentially in a raw state and some basic data processing must be done to convert the data into a usable form.

This document provides a description of this basic data processing, starting with raw lens data, and ending up with time synchronized sensor streams in calibrated physical units.

2 Definitions

2.1 Measurement units

Data is provided from three different sensors: a three axis accelerometer, a three axis gyro (angular rate sensor), and a three axis magnetometer. The raw data from the sensor analog to digital converters (ADC) is output from the lens, and the processing described below will transform this into physically meaningful measurements:

- **Accelerometer**: The accelerometer measures both acceleration and gravity; in this document we refer to this combination as specific force, though it is also sometimes called proper acceleration. Units are m s^{-2}.
- **Gyro**: The gyro measures the angular rate of the system. Units are rad s^{-1}.
- **Magnetometer**: The magnetometer measures the local magnetic field. Units are T (tesla).

All three sensors generate measurements defined in the body frame, as defined in the following section.

2.2 Coordinate frames

2.2.1 Focal plane frame (FP)

The focal plane frame is a right-handed coordinate frame, which has a well-defined physical relationship with the PL mount (the fundamental mechanical interface between the lens and the camera):

- The origin of the focal plane (FP) frame is located at the centre of the camera focal plane.
- The y axis points towards the top of the focal plane.
- The z axis is antiparallel to the direction the camera is looking.
- The x axis completes the right-handed set.

2.2.2 Body frame (B)

All three sensors generate measurements defined in the body frame (abbreviated as the B frame). This is a right-handed frame, with its origin at the physical location of the accelerometer in the lens. The origin and frame orientation relative to the PL mount are both lens-model dependent, but for all currently manufactured prime /i^2 lenses, the definition is as follows:

- The origin of the body (B) frame is at coordinate (0.003, -0.034, -0.08) of the FP frame (units meters). Physically, the is 8 cm in front of the focal plane and 3.4 cm below the optical axis.
- The z axis is antiparallel to the FP frame z axis.
- The y axis is a vector with direction (-0.260188, 0.965558, 0) in the FP frame.
- The x axis completes the right-handed set.

2.2.3 Camera frame (C)

For fusion of video data with inertial data, the effective optical pinhole location for the camera must be defined. This corresponds to the entrance pupil position. Note that the distance between the focal plane and the entrance pupil varies depending on the focus setting and zoom setting of the lens, and that this distance can be determined using the entrance pupil position lens metadata. The definition of the frame is:

- The origin of the camera (C) frame is the centre of the entrance pupil of the lens.
- The axes orientation is identical to the focal plane (FP) frame.

2.2.4 Summary

As an aid to understanding, the relationship between all three coordinate frames is shown below.

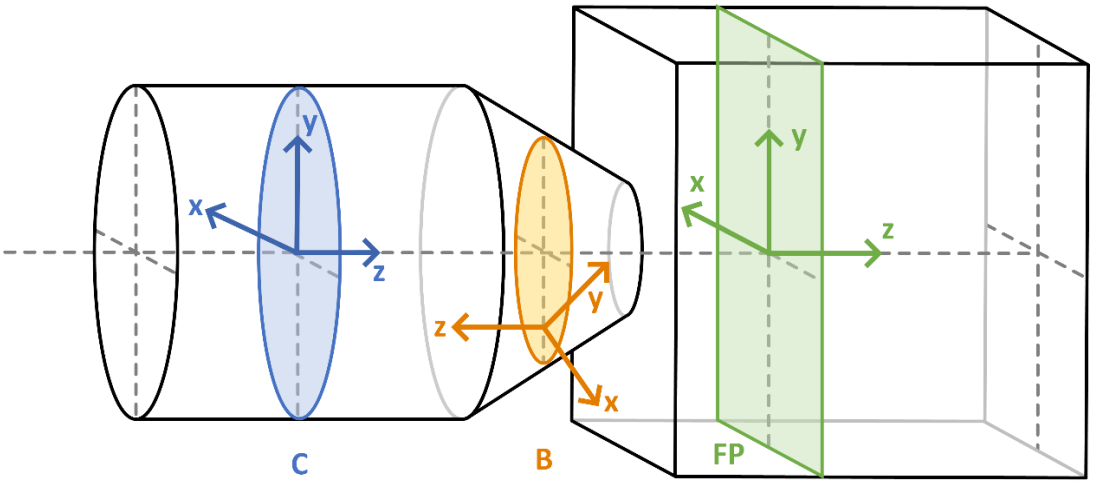

3 Data import

3.1 Main data stream

The first step is to parse the binary Kdi responses. Each Kdi poll will generate a response which contains the following information:

- A sequence number in the range 0-255 which ensures that each Kdi poll from the camera can be unambiguously matched with the response from the lens.
- A high resolution timestamp of when the Kdi poll was received by the lens. This takes the form of a 16 bit unsigned integer, with a clock frequency of 150 kHz.
- A single sample from the magnetometer, taken at the time the Kdi poll was received.
- A set of 0 or more bursts of sensor data, where each burst contains eight samples. Each burst contains:
    - The sensor type (either accelerometer or gyro),
    - A high resolution timestamp for the final sample in the burst,
    - The raw ADC data for the eight samples.

The accelerometer and gyro both output data at roughly 238 Hz. The actual data rate can vary quite significantly from this nominal value, so it is important to use high-resolution timestamps for any analysis which requires high timing accuracy.

The magnetometer is sampled at a much lower frequency (equal to the video frame rate). Accelerometer and gyro data must be integrated to recover useful orientation and position information, and high sample frequency is essential to control numerical errors in this process. In contrast, the

magnetometer provides direct information about orientation, and hence a low output frequency is sufficient.

3.2 Inertial calibration coefficients

The inertial calibration is different on each lens, as it is generated by a calibration procedure during lens manufacture. This can be retrieved from the lens using the K61 command. This command returns 45 single precision floating point numbers in little endian form. These coefficients are divided into three sections:

- The first 12 numbers (48 bytes) form the accelerometer calibration matrix. This is a 3x4 matrix (i.e. three rows, four columns) in column major format (i.e. the first three numbers form the first column, the next three numbers form the second column, and so on).
- The next 21 numbers form the gyro calibration matrix, which is a 3x7 matrix in column major format.
- The last 12 numbers form the magnetometer calibration matrix, which is a 3x4 matrix in column major format.

The reason that the gyro calibration has a larger matrix than the accelerometer or magnetometer is that the gyro calibration includes a correction for sensor non-linearity.

4 Conversion to physical units

Up to this point, all sensor data is in its raw state, i.e. a stream of integers as measured by the ADC. This needs to be converted into calibrated physical units.

Accelerometer data is converted using a matrix multiplication:

$$\begin{pmatrix} a_x \\ a_y \\ a_z \end{pmatrix} = \mathbf{M}_{\text{acc}} \begin{pmatrix} aRaw_x \\ aRaw_y \\ aRaw_z \\ 1 \end{pmatrix}$$

where $aRaw$ are raw accelerometer ADC values, a are calibrated specific force measurements in the body (B) coordinate frame, and $\mathbf{M}_{\text{acc}}$ is the 3x4 accelerometer calibration matrix.

Gyro data is converted by:

$$\begin{pmatrix} \omega_x \\ \omega_y \\ \omega_z \end{pmatrix} = \mathbf{M}_{gyro} \begin{pmatrix} gRaw_x \\ gRaw_y \\ gRaw_z \\ gRaw_x^2 \\ gRaw_y^2 \\ gRaw_z^2 \\ 1 \end{pmatrix}$$

where $gRaw$ are the raw gyro ADC values, ω are calibrated angular rate measurements in the body (B) coordinate frame, and $\mathbf{M}_{\text{gyro}}$ is the 3x7 gyro calibration matrix.

Magnetometer data is converted by:

$$\begin{pmatrix} b_x \\ b_y \\ b_z \end{pmatrix} = \mathbf{M}_{\text{mag}} \begin{pmatrix} bRaw_x \\ bRaw_y \\ bRaw_z \\ 1 \end{pmatrix}$$

where $bRaw$ are raw magnetometer ADC values, b are calibrated magnetic field measurements in the body (B) coordinate frame, and $\mathbf{M}_{\text{mag}}$ is the 3x4 magnetometer calibration matrix.

5 Timing data

The raw timing data for the sensor data is in the form of unsigned 16 bit integers from the high resolution clock, running at 150 kHz. These timestamps are extremely accurate, but there are two immediate problems:

- This clock will rollover roughly every 0.44 seconds, and so doesn't provide an absolute timestamp.
- For the accelerometer and gyro data, only the last sample in each burst has a timestamp.

Some further data processing is therefore needed to obtain usable timestamps for each sample of sensor data.

5.1 Unwrapping Kdi timestamps

Kdi polls should generally be sent once per frame of video data. It is therefore straightforward to directly detect when the clock rolls over in the sequence of Kdi timestamps (because the timestamp will decrease between two consecutive polls), and hence unwrap the timestamps into a continuously increasing sequence.

5.2 Unwrapping burst timestamps

The process for unwrapping burst timestamps is very similar to that used for unwrapping Kdi timestamps. However, it is important to make sure that the two different types of timestamp are unwrapped consistently (i.e. that there isn't a full cycle slip between the two streams). This can be done by checking that the most recent burst for each poll has an age of between 0 and 0.033 seconds.

5.3 Accelerometer and gyro sample timestamps

The burst timestamps specifically refer to the time of the last sample in each burst; the preceding seven samples do not have timestamps in the raw data. These timestamps can be inferred with a high degree of accuracy by using linear interpolation. For the very first burst, timestamps can be inferred via extrapolation, or data from the first burst can be discarded.

6 Example test sequence

For test purposes, it is useful to move the lens through a short sequence of orientations, and confirm that the logged inertial data show the expected signals. As an example, below is some data from the following sequence:

- Start with the FP frame x axis pointing up.
- Rotate the lens by -90 degrees around the FP frame z axis. This will leave the lens with the FP frame y axis pointing up (the normal orientation of the lens when mounted on a camera).
- Rotate the lens by -90 degrees around the FP frame x axis. This will leave the lens with the FP frame z axis pointing up.

- Rotate the lens by -90 degrees around the FP frame y axis. This will leave the lens with the FP frame x axis pointing up.

The resulting accelerometer and gyro signals are shown below:

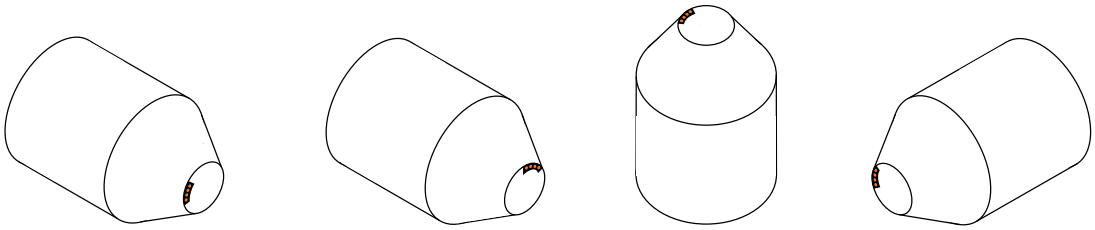

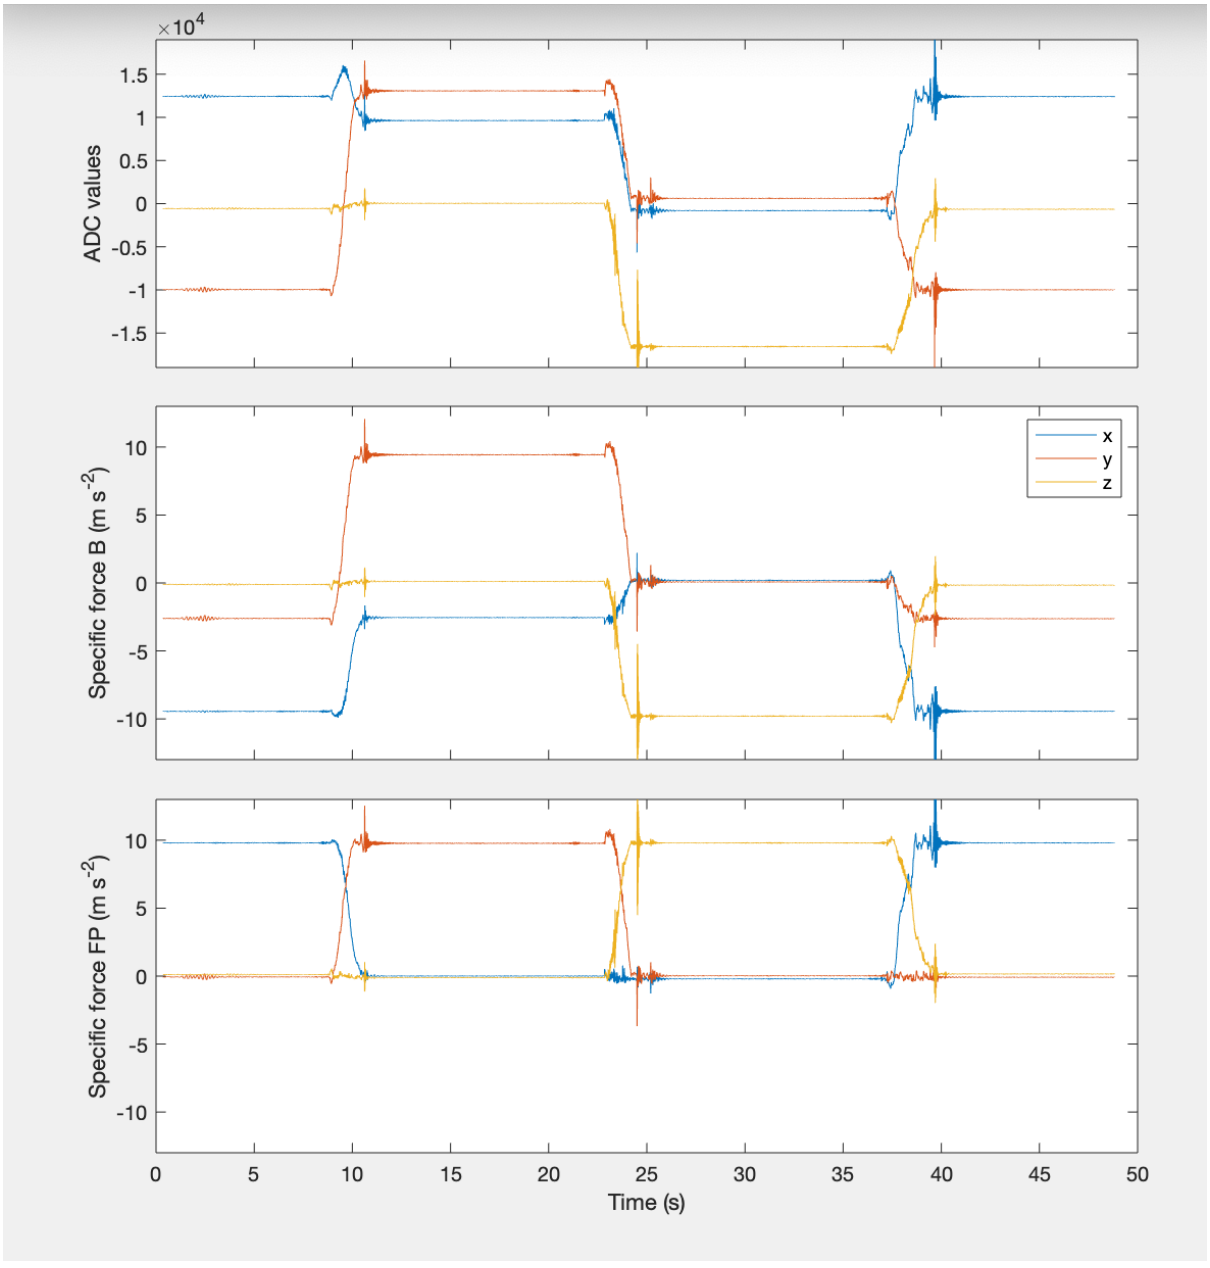
×10^4
ADC values
1.5
1
0.5
0
-0.5
-1
-1.5
Specific force B (m s^-2)
10
5
0
-5
-10
x
y
z
Specific force FP (m s^-2)
10
5
0
-5
-10
0
5
10
15
20
25
30
35
40
45
50
Time (s)

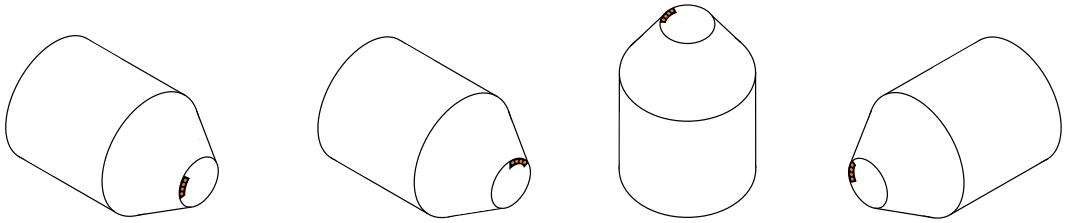

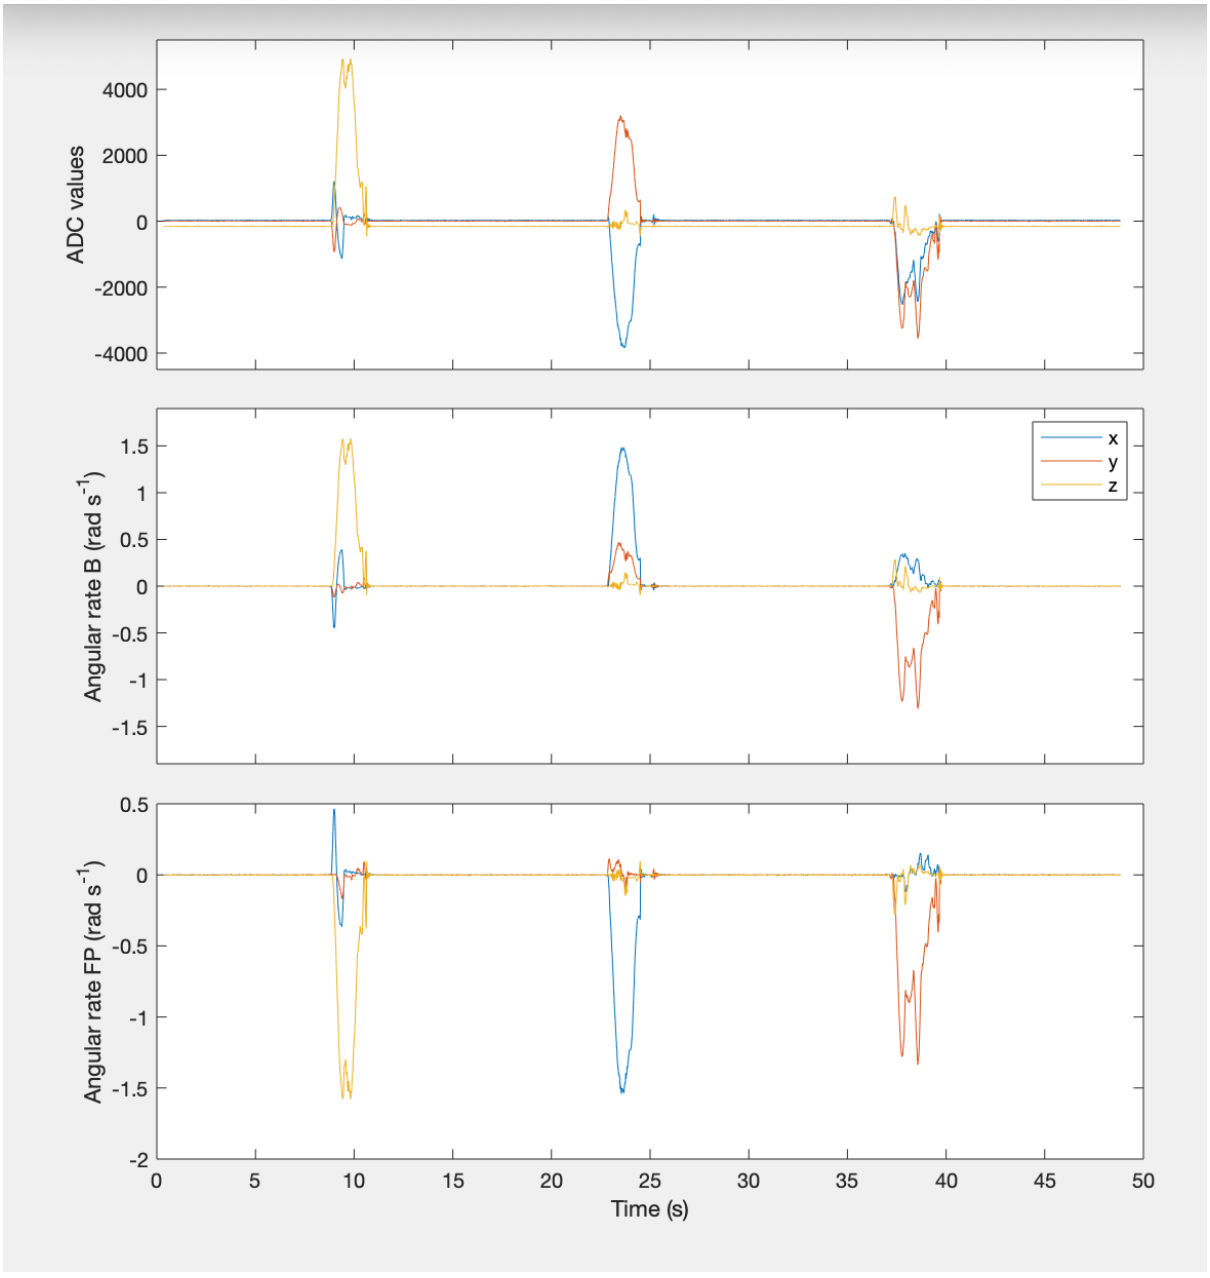
ADC values
4000
2000
0
-2000
-4000
Angular rate B (rad s-1)
1.5
1
0.5
0
-0.5
-1
-1.5
x
y
z
Angular rate FP (rad s-1)
0.5
0
-0.5
-1
-1.5
-2
0
5
10
15
20
25
30
35
40
45
50
Time (s)

Shading Model Design

7 Introduction – Shading Model Design

This document outlines the mathematical model used to represent vignetting in Cooke lenses. Many candidate models were considered before selection of the present model; this model was chosen on the basis that it satisfies the following design criteria:

- Good fit for all lenses, with a worst-case error smaller than 5%.
- Smooth interpolation in between the fitted data points, in both position and iris setting.
- As simple as possible, subject to the above constraints.

8 Model

Vignetting is modelled as a function of the iris setting and position on the detector. The iris setting a is defined as the inverse of the T-stop; this value is proportional to the radius of the physical aperture (the actual constant of proportionality being absorbed by the model parameters). The position parameter x is a distance from the image centre, normalised using the nominal image circle for the lens r_0 (e.g. a point 3 mm from the centre of the sensor would correspond to $x = 0.2$ for a lens with an image circle radius of 15 mm). For anamorphic lenses x is specifically a horizontal offset, and a second position parameter y defines vertical offset; this case both x and y are normalised by the nominal image half-width ($w_0/2$).

8.1 Natural vignetting

Natural vignetting is a continuous variation in the image intensity as a function of position, which occurs even when the lens is stopped down. It is symmetric about $x = 0$, and can be well modelled using a polynomial with only even powers.

$$v_n(x) = 1 - (\alpha_1 x^2 + \alpha_2 x^4 + \alpha_3 x^6), \quad (1)$$

where $v_n(x)$ is the light intensity relative to the intensity at $x = 0$, and α_1, α_2 and α_3 are model coefficients.

8.2 Optical vignetting

Figure 1: Optical vignetting on an example lens, with the iris set to $f/1.4$ (top row) and $f/2$ (bottom row). When viewed straight on (left column), the light collection area is defined purely by the lens's iris. From an angle, the light collection area is defined by the extent of the front and rear elements (at $f/1.4$, top right), or by the front element, rear element, and iris (at $f/2$, bottom right).

Optical vignetting occurs when the edges of optical elements occlude part of the lens's aperture. The effect only occurs at certain combination of iris setting and position on the sensor, so the optical vignetting function is a piecewise function of a and x.

As Figure [1] shows, the occlusion usually comes from more than one element in the lens. In this example, both the front and the rear elements contribute to the effect. For a given a, the onset of vignetting from different elements may occur at different x values, which means that the vignetting function can be composed of several different pieces. The value of the vignetting function is continuous in both x and a, but there is no guarantee that the function derivative will be continuous in either parameter.

The piecewise nature of the function means that simple polynomial models cannot provide acceptable accuracy for most of the test lenses. An explicitly piecewise polynomial or rational function would likely be usable, but such a model would be complicated to specify and evaluate.

The model selected is therefore based on a simple geometric model, which represents the iris, front and rear optical elements as circles. At $x = 0$, all three circles are concentric, and the full area of the aperture circle is available. As x increases, the circles representing the front and rear optical elements

shift horizontally (in opposite directions), and will at some point start to occlude the aperture circle. The optical vignetting effect is then taken as fraction of the iris circle which is unobscured.

In detail, the optical vignetting function is defined as follows. Start by defining the intersection area of three overlapping circles with radii r_1, r_2 and r_3 and with centre distances d_{12} and d_{23} as $B(r_1, r_2, r_3, d_{12}, d_{23})$ (see Section [3] for the full definition). Then, the optical vignetting function $v_o(x, a)$ is defined as

$$v_o(x,a) = \frac{B(a_f(x), a, a_r(x), \beta_1 x, \beta_2 x)}{\pi a^2}, \quad (2)$$

where $a_f(x)$ and $a_r(x)$ are the radii of the front and rear elements respectively, and β_1 and β_2 are model coefficients.

In practice, using fixed constants for a_f and a_r does not provide a sufficiently accurate fit to the true optical vignetting on all lenses, so these values are taken to be an even function of x:

$$a_f(x) = \mu_1 + \mu_2 x^2 + \mu_3 x^4, \quad (3)$$
$$a_r(x) = \nu_1 + \nu_2 x^2 + \nu_3 x^4, \quad (4)$$

where μ_1, μ_2, μ_3 and ν_1, ν_2, ν_3 are model coefficients.

To get the final value for the vignette function $v(x, a)$, the natural and optical vignetting effects are simply multiplied:

$$v(x,a) = v_n(x) v_o(x,a). \quad (5)$$

In total, a spherical lens vignetting model has 11 free parameters: three for the natural vignette, and eight for optical vignette.

8.3 Anamorphic lenses

Anamorphic lenses can also be accommodated by the model outlined above with some minor modifications. Start by defining the radial distance to the image centre: r:

$$r = \sqrt{x^2 + y^2}. \quad (6)$$

The natural vignette can be well modelled using the existing function, in terms of r instead of x:

$$v_n(r) = 1 - (\alpha_1 r^2 + \alpha_2 r^4 + \alpha_3 r^6). \quad (7)$$

On some of the anamorphic lenses, the optical vignette can be handled in a similar way. However, to give good accuracy on all lenses, modified versions Equations 3 and 4 are needed:

$$a_f'(x,y) = \mu_1 + \mu_2 x^2 + \mu_3 y^2 + \mu_4 x^4 + \mu_5 x^2 y^2 + \mu_6 y^4, \quad (8)$$
$$a_r'(x,y) = \nu_1 + \nu_2 x^2 + \nu_3 y^2 + \nu_4 x^4 + \nu_5 x^2 y^2 + \nu_6 y^4, \quad (9)$$

The optical vignetting function can then be reused, with $a_f'(x, y)$ and $a_r'(x, y)$ replacing $a_f(x)$ and $a_r(x)$:

$$v_o(x,y,r,a) = \frac{B(a_f'(x,y), a, a_r'(x,y), \beta_1 r, \beta_2 r)}{\pi a^2}. \quad (10)$$

The two additional six parameters bring the total number of free parameters for anamorphic lenses to 17.

9 Circle intersection area calculation

The optical vignetting model needs an accurate method for calculating the intersection area of three circles. This section outlines an algorithm for this, starting with an expression for the area of two overlapping circles, and then extending this result to cover the full three-circle problem.

9.1 Two-circle intersection area

Define two circles with radii R and r, and with the circle centres separated by a distance d.

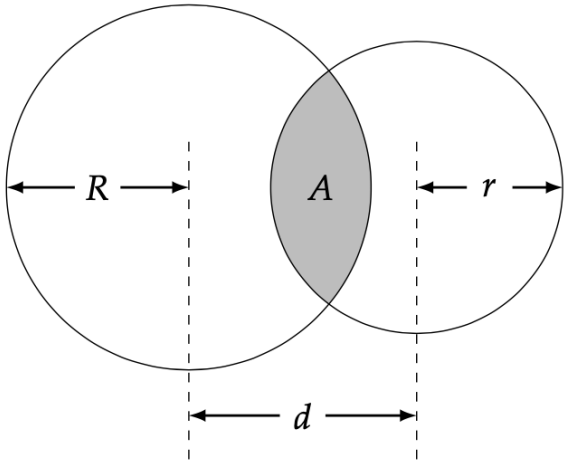

The overlap area $A(R, r, d)$ can be calculated in closed form (see [1] for a derivation):

1. First, check if the area is zero:

$$A = 0 \text{ if } r + R < d \tag{11}$$

If not, then continue with the algorithm.

2. Then, check if one of the circles is nested inside the other circle:

$$A = \begin{cases} \pi r^2 & \text{if } r < R - d \\ \pi R^2 & \text{if } R < r - d \end{cases} \tag{12}$$

If none of these criteria are true, then continue with the algorithm.

3. Finally, the intersecting case has the following solution:

$$A = r^2 \cos^{-1}\left(\frac{d^2 + r^2 - R^2}{2dr}\right) + R^2 \cos^{-1}\left(\frac{d^2 + R^2 - r^2}{2dR}\right) - 1/2\sqrt{(-d + r + R)(d + r - R)(d - r + R)(d + r + R)}. \tag{13}$$

9.2 Three-circle intersection area

Define three circles with collinear centres, circle radii r_1, r_2, r_3, and distance between centres d_{12}, d_{23} and $d_{13} = d_{12} + d_{23}$.

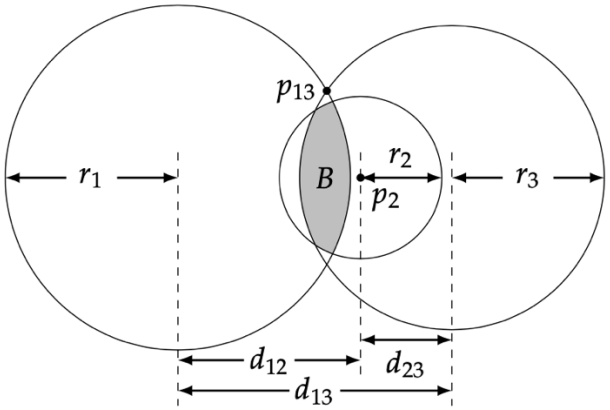

Calculating the intersection area B is more complicated than for the two-circle case, because there are a number of different topologies which must be considered. The following algorithm describes how the area can be found in the general case.

1. First, check if the area is zero:

$$B = \begin{cases} 0 & \text{if } r_1 + r_2 < d_{12} \\ 0 & \text{if } r_2 + r_3 < d_{23} \\ 0 & \text{if } r_1 + r_3 < d_{13}. \end{cases} \tag{14}$$

   If none of these criteria are true, then continue with the algorithm.

2. Then, check if one of the circles is nested inside both of the other circles:

$$B = \begin{cases} \pi r_1^2 & \text{if } r_1 \le r_2 - d_{12} \text{ and } r_1 \le r_3 - d_{13} \\ \pi r_2^2 & \text{if } r_2 \le r_1 - d_{12} \text{ and } r_2 \le r_3 - d_{23} \\ \pi r_3^2 & \text{if } r_3 \le r_1 - d_{13} \text{ and } r_3 \le r_2 - d_{23}. \end{cases} \tag{15}$$

   If none of these criteria are true, then continue with the algorithm.

3. Next, check if one of the circles completely covers another circle. In that case, the problem reduces to that of two overlapping circles, and Section [3.1] gives the result:

$$B = \begin{cases} A(r_2, r_3, d_{23}) & \text{if } r_2 \le r_1 - d_{12} \text{ or } r_3 \le r_1 - d_{13} \\ A(r_1, r_3, d_{13}) & \text{if } r_1 \le r_2 - d_{12} \text{ or } r_3 \le r_2 - d_{23} \\ A(r_1, r_2, d_{12}) & \text{if } r_1 \le r_3 - d_{13} \text{ or } r_2 \le r_3 - d_{23} \end{cases} \tag{16}$$

   If none of these criteria are true, then continue with the algorithm.

4. Finally check for the case where the middle circle is large enough that it does not constrain the intersection area. Calculate the distance $d_{2,13}$ from the centre of the middle circle p_2 to the intersection point between the left and right cirlces p_{13}:

$$d_{2,13} = \sqrt{\frac{d_{12}(r_3^2 - d_{23}^2) + d_{23}(r_1^2 - d_{12}^2)}{d_{12} + d_{23}}}. \quad (17)$$

Then, the area can be determined:

$$B = \begin{cases} A(r_1, r_3, d_{13}) & \text{if } d_{2,13} \le r_2 \\ A(r_1, r_2, d_{12}) + A(r_2, r_3, d_{23}) - \pi r_2^2 & \text{if } d_{2,13} > r_2 \end{cases} \quad (18)$$

10 Data storage format

The model parameters can be downloaded from a lens using the KKi command.

For spherical lenses, coefficients will be returned in the following order:

Index	Coefficient	Suggested variable name	Units
1	α_1	*alpha_1*	
2	α_2	*alpha_2*	
3	α_3	*alpha_3*	
4	β_1	*beta_1*	
5	β_2	*beta_2*	
6	μ_1	*mu_1*	
7	μ_2	*mu_2*	
8	μ_3	*mu_3*	
9	ν_1	*nu_1*	
10	ν_2	*nu_2*	
11	ν_3	*nu_3*	
12	r_0	*nominal_image_radius*	mm

For anamorphic lenses, the coefficients are:

Index	Coefficient	Suggested variable name	Units
1	α_1	*alpha_1*	
2	α_2	*alpha_2*	

3	α_3	*alpha_3*	
4	β_1	*beta_1*	
5	β_2	*beta_2*	
6	μ_1	*mu_1*	
7	μ_2	*mu_2*	
8	μ_3	*mu_3*	
9	μ_4	*mu_4*	
10	μ_5	*mu_5*	
11	μ_6	*mu_6*	
12	ν_1	*nu_1*	
13	ν_2	*nu_2*	
14	ν_3	*nu_3*	
15	ν_4	*nu_4*	
16	ν_5	*nu_5*	
17	ν_6	*nu_6*	
18	w_0	*nominal_image_width*	mm
19	h_0	*nominal_image_height*	mm

11 References

[1] Weisstein, Eric W., “Circle-Circle Intersection.” From *MathWorld–A Wolfram Web Resource.* *http://mathworld.wolfram.com/Circle-CircleIntersection.html*

Spherical Lens Distortion Model

12. Introduction – Spherical Lens Distortion Model

This section defines the mathematical model for lens optical properties, including distortion. The model is intended to support VFX applications (particularly matchmoving), where very high accuracy is required. This aim leads to the following choices:

- Each manufactured lens is measured individually as part of the build process, so model coefficients are as accurate as possible
- The mathematical model accounts for decentering distortion as well as radial distortion
- Model parameters are estimated as a function of the lens focus setting, so an accurate lens model can be reconstructed for each frame of the video

13 Coordinate Frames and Units

The camera frame C is a 3D orthogonal frame:

- The origin is located at the entrance pupil of the lens.
- The y axis points down, and is aligned with the columns of sensor pixels.
- The z axis points in the direction the lens is looking.
- The x axis completes the right-handed set.
- Unit is meters

The pixel frame P is a 2D coordinate system:

- The origin is located at the center of the top-left pixel in the image.
- The x axis points right.
- The y axis points down.

- Unit is pixels.

We also define 2D diagonally normalized coordinates D, which are used for distortion calculations:

- The origin is located at the lens center.
- The x and y axes align with the pixel frame axes.
- Unites are such that the diameter of the lens' nominal image circle is 2.

14 Pinhole Camera Model

The basic pinhole camera model is defined by the following parameters:

- Principle distance f.
- Lens center(c_x, c_y), as an offset from the PL mount center. Positive c_x indicates lens center is shifted to the right; positive c_y indicates lens center is shifted down.

These are both specified in physical units (mm).

15 Distortion Model

For spherical lenses, the Brown-Conradi model is used. For any measured point in the image, the model specifies a correction which will convert the distorted point into an undistorted point. Working in diagonally normalized coordinates, the corrected coordinates are calculated as follows:

$$x' = x(1 + K_1 r^2 + K_2 r^4 + K_3 r^6) + P_1(r^2 + 2x^2) + 2P_2 xy$$
$$y' = x(1 + K_1 r^2 + K_2 r^4 + K_3 r^6) + 2P_1 xy + P_2(r^2 + 2y^2)$$

where:

- x and y specify the measured (i.e. distorted) point, in diagonally normalized coordinates.
- x' and y' specify the corrected (i.e. undistorted) point, in diagonally normalized coordinates.
- r is the distance from the image center, $r^2 = x^2 + y^2$.
- $K_1, K_2$and K_3 are the radial distortion coefficients.
- P_1 and P_2 are the decentering distortion coefficients.

16 Entrance pupil off-axis shift

The basic model described so far assumes that all rays that enter the camera pass through a single point in 3D space, i.e. the pinhole. However, real lenses deviate from this model which can have a significant effect on reconstruction accuracy. The effect is typically important on wide-angle lenses when imaging objects which are relatively close to the camera (roughly 1 meter or less). A more detailed discussion can be seen in reference [1].

Entrance pupil off-axis shift is modelled as a translation of the ray origin along the optical axis, relative to the nominal pinhole location. For small ray angles the shift is defined to be zero, so the paraxial entrance pupil is coincident with the origin of the camera frame (*C*). As the ray angle increases, the ray origin will shift along the *z* axis of the camera frame. The shift can be positive or negative, depending on the lens design. The shift is calculated as follows:

$$z = S_1 r^2 + S_2 r^4$$

where:

- *r* is the distance from the image centre ($r^2 = x^2 + y^2$) of a measured point, in diagonally normalised coordinates (exactly as used in the distortion calculation).
- S_1 and S_2 are the entrance pupil off-axis shift coefficients.
- *z* is the position of the entrance pupil, specified in mm.

17 Model as a Function of Focus Setting

Each of the optical parameters (f, c_x, c_y , K_1, K_2, K_3 , P_1, P_2) is modelled as function of the focus setting. First define the normalized focus setting as:

$$d_n = \frac{s_{min}}{s}$$

where:

- d_n is the normalized focus setting.
- s_{min} is the minimum focus setting for the lens (in mm).
- s is the current focus setting for the lens (in mm).

d_n will fall in the range (0,1) with 0 indicating infinity focus. Lens optical parameters are then defined as a 3rd order polynomial function of d_n. For example, the principle distance f would be evaluated as follows:

$$f = f_0 + f_1 d_n + f_2 d_n^2 + f_3 d_n^3$$

18 Lens Data

The complete set of data forming the optical lens model is stored as 42 single precision floating point numbers.

Line	Values	Description	Units
1	s_{min}	Minimum focus setting	mm
2	a_{nom}	Nominal image circle diameter	mm
3	f_0, f_1, f_2, f_3	Principle distance	mm
4	$c_{x_0}, c_{x_1}, c_{x_2}, c_{x_3}$	Principle point x offset	mm
5	$c_{y_0}, c_{y_1}, c_y, c_{y_3}$	Principle point y offset	mm

6	$K_{1_0}, K_{1_1}, K_{1_2}, K_{1_3}$	Radial distortion K_1	
7	$K_{2_0}, K_{2_1}, K_{2_2}, K_{2_3}$	Radial distortion K_2	
8	$K_{3_0}, K_{3_1}, K_{3_2}, K_{3_3}$	Radial distortion 3	
9	$P_{1_0}, P_{1_1}, P_{1_2}, P_{1_3}$	Decentering distortion P_1	
10	$P_{2_0}, P_{2_1}, P_{2_2}, P_{2_3}$	Decentering distortion P_2	
11	$S_{1_0}, S_{1_1}, S_{1_2}, S_{1_3}$	Entrance pupil off-axis shift S_1	mm
12	$S_{2_0}, S_{2_1}, S_{2_2}, S_{2_3}$	Entrance pupil off-axis shift S_2	mm

References

[1] Gennery, D.B., *"Generalized Camera Calibration Including Fish-Eye Lenses"*, Int J Comput Vision (2006) 68: 239.

CookeOpticsLimited

Cooke Close, Thurmaston
Leicester, LE4 8PT, United Kingdom

T +44 (0) 116 264 0700
F +44 (0) 116 264 0707
E lenses@cookeoptics.com
W cookeoptics.com